User Guide for Cisco Security Manager 4.18

First Published: July, 2018

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco website at www.cisco.com/go/offices.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

Any Internet Protocol (IP) addresses used in this document are not intended to be actual addresses. Any examples, command display output, and figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses in illustrative content is unintentional and coincidental.

User Guide for Cisco Security Manager 4.18
© 2018 Cisco Systems, Inc. All rights reserved.

Preface lxiii

Conventions lxiii

Obtain Documentation and Submit a Service Request lxiv

The Basics of Using Security Manager

CHAPTER 1

Getting Started with Security Manager 1-1

Product Overview 1-1

Primary Benefits of Cisco Security Manager 1-2

Security Manager Policy Feature Sets 1-4

Security Manager Applications Overview 1-6

Device Monitoring Overview 1-7

IPv6 Support in Security Manager 1-8

Configuring IPv6 on Security Manager Server 1-9

Configuring IPv6 Policies 1-9

Policy Object Changes in Security Manager 4.4 1-10

Logging In to and Exiting Security Manager 1-11

Understanding User Permissions 1-11

Logging In to the Cisco Security Management Suite Server 1-12

Logging In to and Exiting the Security Manager Client 1-12

Using Configuration Manager - Overview 1-14

Configuration Manager Overview 1-14

Device View Overview 1-15

Policy View Overview 1-16

Map View Overview 1-18

Task Flow for Configuring Security Policies 1-19

Policy and Policy Object Overview 1-20

Workflow and Activities Overview 1-20

Working in Workflow Mode 1-21

Working in Non-Workflow Mode 1-22

Comparing Workflow Modes 1-22

Using the JumpStart to Learn About Security Manager 1-24

Completing the Initial Security Manager Configuration 1-25

Configuring an SMTP Server and Default Addresses for E-Mail Notifications 1-27
Changing Workflow Modes 1-28
Understanding Basic Security Manager Interface Features 1-29
Menu Bar Reference for Configuration Manager 1-29
File Menu (Configuration Manager) 1-30
Edit Menu (Configuration Manager) 1-31
View Menu (Configuration Manager) 1-31
Policy Menu (Configuration Manager) 1-32
Map Menu (Configuration Manager) 1-33
Manage Menu (Configuration Manager) 1-34
Tools Menu (Configuration Manager) 1-34
Activities Menu (Configuration Manager) 1-36
Tickets Menu (Configuration Manager) 1-36
Launch Menu (Configuration Manager) 1-37
Help Menu (Configuration Manager) 1-38
Toolbar Reference (Configuration Manager) 1-39
Using Global Search 1-42
Using Selectors 1-45
Filtering Items in Selectors 1-45
Create Filter Dialog Box 1-46
Using Wizards 1-47
Using Tables 1-48
Filtering Tables 1-48
Table Columns and Column Heading Features 1-49
Using Text Fields 1-49
Understanding ASCII Limitations for Text 1-50
Finding Text in Text Boxes 1-50
Navigating Within Text Boxes 1-50
Selecting or Specifying a File or Directory in Security Manager 1-50
Troubleshooting User Interface Problems 1-51
Accessing Online Help 1-52

CHAPTER 2 **Preparing Devices for Management** 2-1
Understanding Device Communication Requirements 2-1
Setting Up SSL (HTTPS) 2-3
Setting Up SSL (HTTPS) on PIX Firewall, ASA and FWSM Devices 2-3
Setting Up SSL on Cisco IOS Routers 2-4
Setting Up SSH 2-5
Critical Line-Ending Conventions for SSH 2-5

Testing Authentication 2-6
Setting Up SSH on Cisco IOS Routers, Catalyst Switches, and Catalyst 6500/7600 devices 2-6
Preventing Non-SSH Connections (Optional) 2-7
Setting Up AUS or Configuration Engine 2-8
Setting Up AUS on PIX Firewall and ASA Devices 2-8
Configuring Licenses on Cisco ASA Devices 2-9
Configuring Licenses on Cisco IOS Devices 2-10
Initializing IPS Devices 2-10

CHAPTER 3 **Managing the Device Inventory** 3-1
Understanding the Device Inventory 3-1
Understanding the Device View 3-1
Understanding Device Names and What Is Considered a Device 3-3
Understanding Device Credentials 3-4
Understanding Device Properties 3-6
Adding Devices to the Device Inventory 3-6
Working with Generically Supported Devices 3-8
Working with Device Clusters 3-9
Adding Devices from the Network 3-12
Device Information Page – Add Device from Network 3-14
Service Module Credentials Dialog Box 3-19
IPS Module Discovery Dialog Box 3-20
Adding Devices from Configuration Files 3-22
Device Information Page—Configuration File 3-23
Adding Devices by Manual Definition 3-26
Device Information Page—New Device 3-27
Adding Devices from an Inventory File 3-31
Device Information Page—Add Device from File 3-33
Working with the Device Inventory 3-36
Adding, Editing, or Deleting Auto Update Servers or Configuration Engines 3-36
Server Properties Dialog Box 3-38
Available Servers Dialog Box 3-39
Adding or Changing Interface Modules 3-40
Viewing or Changing Device Properties 3-40
Device Properties: General Page 3-41
Device Credentials Page 3-45
Device Groups Page 3-49
Cluster Information Page 3-50
Policy Object Override Pages 3-52

Changing Critical Device Properties 3-52
Image Version Changes That Do Not Change the Feature Set in Security Manager 3-53
Changes That Change the Feature Set in Security Manager 3-54
Showing Device Containment 3-56
Cloning a Device 3-56
Deleting Devices from the Security Manager Inventory 3-58
Device Delete Validation Dialog Box 3-59
Working with Device Groups 3-59
Understanding Device Grouping 3-60
Edit Device Groups Dialog Box 3-61
Creating Device Group Types 3-62
Creating Device Groups 3-62
Deleting Device Groups or Group Types 3-63
Adding Devices to or Removing Them From Device Groups 3-63
Working with Device Status View 3-64

CHAPTER 4 **Managing Activities** 4-1
Understanding Activities 4-1
Benefits of Activities 4-2
Activity Approval 4-3
Activities and Locking 4-3
Activities and Multiple Users 4-4
Understanding Activity/Ticket States 4-4
Working with Activities/Tickets 4-7
Accessing Activity Functions in Workflow Mode 4-8
Accessing Ticket Functions in Non-Workflow Mode 4-9
Activity/Ticket Manager Window 4-10
Creating an Activity/Ticket 4-14
Responding to the Activity/Ticket Required Dialog Box 4-14
Opening an Activity/Ticket 4-15
Closing an Activity/Ticket 4-16
Viewing Change Reports 4-16
Selecting a Change Report in Non-Workflow Mode with Ticket Management Disabled 4-18
Validating an Activity/Ticket 4-18
Submitting an Activity for Approval (Workflow Mode with Activity Approver) 4-20
Approving or Rejecting an Activity (Workflow Mode) 4-21
Discarding an Activity/Ticket 4-22
Viewing Activity/Ticket Status and History 4-23

CHAPTER 5 **Managing Policies** 5-1

Understanding Policies 5-1
Settings-Based Policies vs. Rule-Based Policies 5-2
Service Policies vs. Platform-Specific Policies 5-2
Local Policies vs. Shared Policies 5-3
Understanding Rule Inheritance 5-4
Inheritance vs. Assignment 5-6
Policy Management and Objects 5-7
Understanding Policy Locking 5-8
Understanding Locking and Policies 5-9
Understanding Locking and VPN Topologies 5-10
Understanding Locking and Objects 5-10
Customizing Policy Management for Routers and Firewall Devices 5-11
Discovering Policies 5-12
Discovering Policies on Devices Already in Security Manager 5-15
Create Discovery Task and Bulk Rediscovery Dialog Boxes 5-18
Viewing Policy Discovery Task Status 5-22
Discovery Status Dialog Box 5-23
Policy Discovery Status Page 5-25
Frequently Asked Questions about Policy Discovery 5-27
Managing Policies in Device View and the Site-to-Site VPN Manager 5-30
Policy Status Icons 5-30
Performing Basic Policy Management 5-31
Configuring Local Policies in Device View 5-31
Copying Policies Between Devices 5-33
Unassigning a Policy 5-36
Working with Shared Policies in Device View or the Site-to-Site VPN Manager 5-37
Using the Policy Banner 5-38
Policy Shortcut Menu Commands in Device View and the Site-to-Site VPN Manager 5-40
Sharing a Local Policy 5-41
Sharing Multiple Policies of a Selected Device 5-42
Unsharing a Policy 5-43
Assigning a Shared Policy to a Device or VPN Topology 5-44
Adding Local Rules to a Shared Policy 5-45
Inheriting or Uninheriting Rules 5-47
Cloning (Copying) a Shared Policy 5-47
Renaming a Shared Policy 5-48
Modifying Shared Policy Definitions in Device View or the Site-to-Site VPN Manager 5-49
Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager 5-49

Managing Shared Policies in Policy View 5-50
Policy View Selectors 5-52
Policy View—Shared Policy Selector Options 5-53
Creating a New Shared Policy 5-54
Modifying Policy Assignments in Policy View 5-54
Deleting a Shared Policy 5-56
Managing Policy Bundles 5-57
Creating a New Policy Bundle 5-57
Cloning a Policy Bundle 5-58
Renaming a Policy Bundle 5-58
Assigning Policy Bundles to Devices 5-59

CHAPTER 6 **Managing Policy Objects** 6-1
Selecting Objects for Policies 6-2
Policy Object Manager 6-4
Policy Object Manager: Undocking and Docking 6-8
Policy Object Manager Shortcut Menu 6-8
Working with Policy Objects—Basic Procedures 6-9
Creating Policy Objects 6-9
Editing Objects 6-12
Using Category Objects 6-13
Cloning (Duplicating) Objects 6-14
Viewing Object Details 6-14
Generating Object Usage Reports 6-15
Deleting Objects 6-16
Managing Object Overrides 6-17
Understanding Policy Object Overrides for Individual Devices 6-18
Allowing a Policy Object to Be Overridden 6-18
Creating or Editing Object Overrides for a Single Device 6-19
Creating or Editing Object Overrides for Multiple Devices At A Time 6-19
Deleting Device-Level Object Overrides 6-21
Overridable Objects in Security Manager 6-22
Importing and Exporting Policy Objects 6-23
Understanding AAA Server and Server Group Objects 6-27
Supported AAA Server Types 6-28
Additional AAA Support on ASA, PIX, and FWSM Devices 6-28
Predefined AAA Authentication Server Groups 6-30
Default AAA Server Groups and IOS Devices 6-31
Creating AAA Server Objects 6-32

Add or Edit AAA Server Dialog Box 6-33
AAA Server Dialog Box—RADIUS Settings 6-35
AAA Server Dialog Box—TACACS+ Settings 6-38
AAA Server Dialog Box—Kerberos Settings 6-39
AAA Server Dialog Box—LDAP Settings 6-40
AAA Server Dialog Box—NT Settings 6-43
AAA Server Dialog Box—SDI Settings 6-43
AAA Server Dialog Box—HTTP-FORM Settings 6-44
Add and Edit LDAP Attribute Map Dialog Boxes 6-46
Add and Edit LDAP Attribute Map Value Dialog Boxes 6-47
Add and Edit Map Value Dialog Boxes 6-47
Creating AAA Server Group Objects 6-48
AAA Server Group Dialog Box 6-49
Creating Access Control List Objects 6-53
Creating Extended Access Control List Objects 6-54
Creating Standard Access Control List Objects 6-56
Creating Web Access Control List Objects 6-57
Creating Unified Access Control List Objects 6-58
Add or Edit Access List Dialog Boxes 6-59
Add and Edit Extended Access Control Entry Dialog Boxes 6-61
Add and Edit Standard Access Control Entry Dialog Boxes 6-64
Add and Edit Web Access Control Entry Dialog Boxes 6-65
Add and Edit Unified Access Control Entry Dialog Boxes 6-67
Configuring Time Range Objects 6-71
Recurring Ranges Dialog Box 6-72
Understanding Interface Role Objects 6-73
Creating Interface Role Objects 6-74
Interface Role Dialog Box 6-75
Specifying Interfaces During Policy Definition 6-76
Using Interface Roles When a Single Interface Specification is Allowed 6-77
Handling Name Conflicts between Interfaces and Interface Roles 6-78
Understanding Map Objects 6-78
Understanding Networks/Hosts Objects 6-80
Contiguous and Discontiguous Network Masks for IPv4 Addresses 6-81
Creating Networks/Hosts Objects 6-82
Add or Edit Network/Host Dialog Box 6-83
Using Unspecified Networks/Hosts Objects 6-86
Specifying IP Addresses During Policy Definition 6-87
VM Attribute Policies 6-89

Communication between the VM attribute agent and vCenter 6-89
Attribute Agent States 6-90
Guidelines for Configuring vCenter Virtual Machines 6-90
Configuring VM Attribute Policies 6-91
Understanding Pool Objects 6-92
Add or Edit IPv4 Pool Dialog Box 6-92
Add or Edit IPv6 Pool Dialog Box 6-93
Add or Edit MAC Address Pool Dialog Box 6-94
Add or Edit NET Pool Object Dialog Box 6-95
Add or Edit DHCPv6 Pool Dialog Box 6-96
Configuring SAML Identity Provider 6-98
Adding or Editing SAML Identity Provider 6-98
Understanding and Specifying Services and Service and Port List Objects 6-100
Configuring Port List Objects 6-102
Configuring Service Objects 6-103
How Policy Objects are Provisioned as Object Groups 6-106
How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups 6-107
How Service Objects are Provisioned as Object Groups 6-108

CHAPTER 7 **Managing FlexConfigs** 7-1

Understanding FlexConfig Policies and Policy Objects 7-2
Using CLI Commands in FlexConfig Policy Objects 7-2
Using Scripting Language Instructions 7-3
Scripting Language Example 1: Looping 7-3
Scripting Language Example 2: Looping with Two-Dimensional Arrays 7-3
Example 3: Looping with If/Else Statements 7-4
Understanding FlexConfig Object Variables 7-5
Example of FlexConfig Policy Object Variables 7-6
FlexConfig System Variables 7-7
Predefined FlexConfig Policy Objects 7-19
Configuring FlexConfig Policies and Policy Objects 7-25
A FlexConfig Creation Scenario 7-25
Creating FlexConfig Policy Objects 7-28
Add or Edit FlexConfig Dialog Box 7-30
Create Text Object Dialog Box 7-32
Add or Edit Text Object Dialog Box 7-32
FlexConfig Undefined Variables Dialog Box 7-33
Property Selector Dialog Box 7-34

Editing FlexConfig Policies 7-35
FlexConfig Policy Page 7-36
Values Assignment Dialog Box 7-37
FlexConfig Preview Dialog Box 7-38
Troubleshooting FlexConfigs 7-38

CHAPTER 8 **Managing Deployment** 8-1
Understanding Deployment 8-1
Overview of the Deployment Process 8-1
Deployment in Non-Workflow Mode 8-3
Deployment Task Flow in Non-Workflow Mode 8-3
Job States in Non-Workflow Mode 8-4
Deployment in Workflow Mode 8-5
Deployment Task Flow in Workflow Mode 8-5
Job States in Workflow Mode 8-6
Deployment Job Approval 8-7
Deployment Jobs and Multiple Users 8-8
Including Devices in Deployment Jobs or Schedules 8-8
Understanding Deployment Methods 8-8
Deploying Directly to a Device 8-9
Deploying to a Device through an Intermediate Server 8-10
Deploying to a File 8-11
Understanding How Out-of-Band Changes are Handled 8-12
Handling Device OS Version Mismatches 8-13
Overview of the Deployment Manager and Configuration Archive 8-15
Understanding What You Can Do with the Deployment Manager 8-16
Deployment Manager Window 8-16
Deployment Workflow Commentary Dialog Box 8-20
Deployment Schedules Tab, Deployment Manager 8-21
Configuration Archive Window 8-23
Working with Deployment and the Configuration Archive 8-25
Viewing Deployment Status and History for Jobs and Schedules 8-26
Tips for Successful Deployment Jobs 8-27
Deploying Configurations in Non-Workflow Mode 8-28
Edit Deploy Method Dialog Box 8-30
Warning - Partial VPN Deployment Dialog Box 8-31
Deployment Status Details Dialog Box 8-32
Deploying Configurations in Workflow Mode 8-34
Creating and Editing Deployment Jobs 8-35

Submitting Deployment Jobs 8-38
Approving and Rejecting Deployment Jobs 8-39
Deploying a Deployment Job in Workflow Mode 8-39
Discarding Deployment Jobs 8-41
Deploying Configurations Using an Auto Update Server or CNS Configuration Engine 8-41
Deploying Configurations to a Token Management Server 8-43
Previewing Configurations 8-44
Detecting and Analyzing Out of Band Changes 8-45
Exceptions to Out of Band Change Detection 8-47
OOB (Out of Band) Changes Dialog Box 8-48
OOB Re-sync. Tool 8-50
Redeploying Configurations to Devices 8-53
Aborting Deployment Jobs 8-55
Creating or Editing Deployment Schedules 8-55
Schedule Dialog Box 8-56
Add Other Devices Dialog Box 8-58
Suspending or Resuming Deployment Schedules 8-58
Adding Configuration Versions from a Device to the Configuration Archive 8-59
Viewing and Comparing Archived Configuration Versions 8-59
Configuration Version Viewer 8-60
Viewing Deployment Transcripts 8-62
Rolling Back Configurations 8-63
Understanding Configuration Rollback 8-63
Understanding Rollback for Devices in Multiple Context Mode 8-64
Understanding Rollback for Failover Devices 8-65
Understanding Rollback for Catalyst 6500/7600 Devices 8-65
Understanding Rollback for IPS and IOS IPS 8-66
Commands that Can Cause Conflicts after Rollback 8-67
Commands to Recover from Failover Misconfiguration after Rollback 8-68
Rolling Back Configurations to Devices Using the Deployment Manager 8-69
Using Rollback to Deploy Archived Configurations 8-70
Performing Rollback When Deploying to a File 8-71

CHAPTER 9 **Troubleshooting Device Communication and Deployment** 9-1
Testing Device Connectivity 9-1
Device Connectivity Test Dialog Box 9-3
Managing Device Communication Settings and Certificates 9-4
Multiple Certificate Authentication Support 9-4
Manually Adding SSL Certificates for Devices that Use HTTPS Communications 9-5

Security Certificate Rejected When Discovering Device 9-6
Invalid Certificate Error During Device Discovery 9-7
Troubleshooting SSH Connection Problems 9-7
Troubleshooting Device Communication Failures 9-8
Resolving Red X Marks in the Device Selector 9-9
Troubleshooting Deployment 9-9
Changing How Security Manager Responds to Device Messages 9-10
Memory Violation Deployment Errors for ASA 8.3+ Devices 9-12
Error While Attempting to Remove Unreferenced Object 9-12
Security Manager Unable to Communicate With Device After Deployment 9-12
Updating VPNs That Include Routing Processes 9-13
Mixing Deployment Methods with Router and VPN Policies 9-14
Deployment Failures for Routers 9-14
Deployment Failures for Catalyst Switches and Service Modules 9-16
Changing How Security Manager Deploys Configurations to Multiple-Context FWSM 9-17
Deployment Failures to Devices Managed by AUS 9-18
Troubleshooting the Setup of Configuration Engine-Managed Devices 9-18

CHAPTER 10 **Managing the Security Manager Server** 10-1
Overview of Security Manager Server Management and Administration 10-1
Managing a Cluster of Security Manager Servers 10-2
Overview of Security Manager Server Cluster Management 10-2
Splitting a Security Manager Server 10-3
Synchronizing Shared Policies Among Security Manager Servers 10-5
Exporting the Device Inventory 10-6
Exporting the Device Inventory from the Security Manager Client 10-6
Supported CSV Formats for Inventory Import/Export 10-9
Exporting the Device Inventory from the Command Line 10-10
Exporting Shared Policies 10-12
Importing Policies or Devices 10-13
Installing Security Manager License Files 10-16
Certificate Trust Management 10-18
Working with Audit Reports 10-19
Understanding Audit Reports 10-19
Generating the Audit Report 10-20
Using the Audit Report Window 10-21
Purging Audit Log Entries 10-23
Taking Over Another User's Work 10-23
Changing Passwords for the Admin or Other Users 10-24

Backing up and Restoring the Security Manager Database 10-24
Backing Up the Server Database 10-25
Restoring the Server Database 10-27
Generating Data for the Cisco Technical Assistance Center 10-28
Creating Diagnostics Files for the Cisco Technical Assistance Center 10-28
Generating Deployment or Discovery Status Reports 10-30
Generating a Partial Database Backup for the Cisco Technical Assistance Center 10-30

CHAPTER 11 **Configuring Security Manager Administrative Settings** 11-1
API Settings Page 11-2
AutoLink Settings Page 11-3
ACL Hit Count Settings Page 11-3
CCO Settings Page 11-4
Configuration Archive Page 11-6
CS-MARS Page 11-7
New or Edit CS-MARS Device Dialog Box 11-8
CSM Mobile Page 11-9
Customize Desktop Page 11-10
Debug Options Page 11-11
Deployment Page 11-13
Device Communication Page 11-21
Add Certificate Dialog Box 11-24
Device Groups Page 11-24
Discovery Page 11-25
Event Management Page 11-27
Troubleshooting Syslog Relay Servers 11-33
Device Management via IP 11-33
CPU Throttling Policy Dialog Box 11-33
Syslog Relay Statistics Dialog Box 11-35
Health and Performance Monitor Page 11-36
Report Manager Page 11-38
Identity Settings Page 11-38
Image Manager Page 11-41
IP Intelligence Settings Page 11-41
Eventing Notification Settings Page 11-45
IPS Updates Page 11-47
Edit Update Server Settings Dialog Box 11-52

Edit Auto Update Settings Dialog Box 11-54
Edit Signature Download Filter Settings Dialog Box 11-55
ISE Settings Page 11-56
Licensing Page 11-57
CSM Tab, Licensing Page 11-57
IPS Tab, Licensing Page 11-58
Verifying IPS Devices for License Update or Redeployment 11-60
Selecting IPS License Files 11-61
License Update Status Details Dialog Box 11-62
Logs Page 11-62
Policy Management Page 11-64
Policy Objects Page 11-66
Process Monitoring Settings Page 11-67
Single Sign-on Configuration Page 11-68
Rule Expiration Page 11-69
Server Security Page 11-70
Take Over User Session Page 11-71
Ticket Management Page 11-72
Token Management Page 11-73
VPN Policy Defaults Page 11-74
Workflow Page 11-75
Wall Settings Page 11-77

Firewall Services and NAT

CHAPTER 12 **Introduction to Firewall Services** 12-1
Overview of Firewall Services 12-1
Understanding the Processing Order of Firewall Rules 12-2
Understanding How NAT Affects Firewall Rules 12-3
ACL Names Preserved by Security Manager 12-4
ACL Naming Conventions 12-5
Resolving User Defined ACL Policy Naming Conflicts 12-6
Resolving ACL Name Conflicts Between Policies 12-7
Managing Your Rules Tables 12-7
Using Rules Tables 12-8
Adding and Removing Rules 12-9
Editing Rules 12-10

Adding or Editing Address Cells in Rules Tables 12-11
Adding or Editing User Cells in Rules Tables 12-12
Adding or Editing Services Cells in Rules Tables 12-13
Adding or Editing Interfaces or Zones Cells in Rules Tables 12-13
Editing Category Cells in Rules Tables 12-14
Editing Description Cells in Rules Tables 12-14
Showing the Contents of Cells in Rules Tables 12-14
Finding and Replacing Items in Rules Tables 12-16
Find and Replace Dialog Box 12-17
Moving Rules and the Importance of Rule Order 12-19
Enabling and Disabling Rules 12-20
Using Sections to Organize Rules Tables 12-20
Add and Edit Rule Section Dialog Boxes 12-22
Combining Rules 12-22
Combine Rules Selection Summary Dialog Box 12-24
Interpreting Rule Combiner Results 12-25
Example Rule Combiner Results 12-27
Converting IPv4 Rules to Unified Rules 12-28
Generating Policy Query Reports 12-28
Querying Device or Policy Dialog Box 12-29
Interpreting Policy Query Results 12-32
Example Policy Query Result 12-34
Optimizing Network Object Groups When Deploying Firewall Rules 12-35
Expanding Object Groups During Discovery 12-35

CHAPTER 13 **Managing Identity-Aware Firewall Policies** 13-1
Overview of Identity-Aware Firewall Policies 13-1
User Identity Acquisition 13-2
Requirements for Identity-Aware Firewall Policies 13-3
Configuring the Firewall to Provide Identity-Aware Services 13-7
Configuring Identity-Aware Firewall Policies 13-7
Enabling Identity-Aware Firewall Services 13-8
Identifying Active Directory Servers and Agents 13-8
Configuring Identity Options 13-15
Creating Identity User Group Objects 13-19
Selecting Identity Users in Policies 13-21
Configuring Identity-Based Firewall Rules 13-21
Configuring Cut-Through Proxy 13-23
Collecting User Statistics 13-25

Filtering VPN Traffic with Identity-Based Rules 13-26
Monitoring Identity Firewall Policies 13-27

CHAPTER 14 **Managing TrustSec Firewall Policies 14-1**
Overview of TrustSec Firewall Policies 14-1
Understanding SGT and SXP Support in Cisco TrustSec 14-2
Roles in the Cisco TrustSec Solution 14-3
Security Group Policy Enforcement 14-3
About Speaker and Listener Roles 14-6
Prerequisites for Integrating an ASA with Cisco TrustSec 14-6
Configuring TrustSec Firewall Policies 14-7
Configuring Cisco TrustSec Services 14-8
Configuring Security Exchange Protocol (SXP) Settings 14-8
Defining SXP Connection Peers 14-12
Creating Security Group Objects 14-14
Selecting Security Groups in Policies 14-16
Configuring TrustSec-Based Firewall Rules 14-17
Monitoring TrustSec Firewall Policies 14-17

CHAPTER 15 **Managing Firewall AAA Rules 15-1**
Understanding AAA Rules 15-1
Understanding How Users Authenticate 15-2
Configuring AAA Rules for ASA, PIX, and FWSM Devices 15-4
Configuring AAA Rules for IOS Devices 15-7
AAA Rules Page 15-10
Add and Edit AAA Rule Dialog Boxes 15-13
Edit AAA Option Dialog Box 15-19
AuthProxy Dialog Box 15-19
Edit Server Group Dialog Box 15-19
AAA Firewall Settings Policies 15-20
AAA Firewall Settings Page, Advanced Setting Tab 15-20
Interactive Authentication Configuration Dialog Box 15-24
Clear Connection Configuration Dialog Box 15-25
AAA Firewall Page, MAC-Exempt List Tab 15-26
Firewall AAA MAC Exempt Setting Dialog Box 15-27
AAA Page 15-28
Firewall AAA IOS Timeout Value Setting 15-30

CHAPTER 16 **Managing Firewall Access Rules** 16-1

Understanding Access Rules 16-1
Understanding Global Access Rules 16-3
Understanding Device Specific Access Rule Behavior 16-4
Understanding Access Rule Address Requirements and How Rules Are Deployed 16-5
Configuring Access Rules 16-7
Access Rules Page 16-10
Add and Edit Access Rule Dialog Boxes 16-14
Advanced and Edit Options Dialog Boxes 16-17
Hit Count Selection Summary Dialog Box 16-20
Configuring Expiration Dates for Access Rules 16-22
Configuring Settings for Access Control 16-23
Access Control Settings Page 16-24
Firewall ACL Setting Dialog Box 16-26
Using Automatic Conflict Detection 16-28
Understanding Automatic Conflict Detection 16-28
Understanding the Automatic Conflict Detection User Interface 16-30
Resolving Conflicts 16-34
Viewing Hit Count Details 16-36
Sample Hit Count Details Window 16-38
Importing Rules 16-40
Import Rules Wizard—Enter Parameters Page 16-41
Import Rules Wizard—Status Page 16-42
Import Rules Wizard—Preview Page 16-43
Examples of Imported Rules 16-44
Optimizing Access Rules Automatically During Deployment 16-46
Customizing defaults in the Add Access Rule dialog 16-48

CHAPTER 17 **Managing Firewall Inspection Rules** 17-1

Understanding Inspection Rules 17-2
Choosing the Interfaces for Inspection Rules 17-2
Selecting Which Protocols To Inspect 17-3
Understanding Access Rule Requirements for Inspection Rules 17-4
Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices 17-5
Configuring Inspection Rules 17-5
Inspection Rules Page 17-8
Add or Edit Inspect/Application FW Rule Wizard 17-11
Add or Edit Inspect/Application FW Rule Wizard, Step 2 17-13

Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page 17-17
Configure DNS Dialog Box 17-19
Configure SMTP Dialog Box 17-20
Configure ESMTP Dialog Box 17-20
Configure Fragments Dialog Box 17-20
Configure IMAP or POP3 Dialog Boxes 17-21
Configure RPC Dialog Box 17-21
Custom Protocol Dialog Box 17-22
Configure Dialog Box 17-22
Configuring Protocols and Maps for Inspection 17-22
Configuring Class Maps for Inspection Policies 17-28
Configuring DCE/RPC Maps 17-29
DCE/RPC Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-31
Configuring DNS Maps 17-32
DNS Map Protocol Conformance Tab 17-33
DNS Map Filtering Tab 17-34
DNS Umbrella Connector Tab 17-35
DNS Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-36
Configuring ESMTP Maps 17-39
ESMTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-40
Configuring FTP Maps 17-42
FTP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-43
Configuring GTP Maps 17-45
Add and Edit Country Network Codes Dialog Boxes 17-48
Add and Edit Permit Response Dialog Boxes 17-48
GTP Map Timeouts Dialog Box 17-48
GTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-49
Configuring H.323 Maps 17-51
Add or Edit HSI Group Dialog Boxes 17-53
Add or Edit HSI Endpoint IP Address Dialog Boxes 17-54
H.323 Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-54
Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices 17-56
HTTP Map General Tab 17-57
HTTP Map Entity Length Tab 17-58
HTTP Map RFC Request Method Tab 17-60
HTTP Map Extension Request Method Tab 17-61
HTTP Map Port Misuse Tab 17-62
HTTP Map Transfer Encoding Tab 17-63
Configuring HTTP Maps for ASA 7.2+ and PIX 7.2+ Devices 17-64

HTTP Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes 17-66
Configuring IM Maps for ASA 7.2+, PIX 7.2+ Devices 17-70
IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes 17-71
Configuring IM Maps for IOS Devices 17-73
Configuring IP Options Maps 17-75
Configuring IPv6 Maps 17-77
IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-78
Configuring IPsec Pass Through Maps 17-80
Configuring NetBIOS Maps 17-81
Configuring ScanSafe Maps 17-82
Configuring SIP Maps 17-83
SIP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-85
Configuring Skinny Maps 17-87
Skinny Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-89
Configuring SNMP Maps 17-90
Configuring SCTP Maps 17-91
SCTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-92
Configuring Diameter Maps 17-93
Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes 17-95
Create and Add Custom AVPs 17-97
Create and Add TLS Proxy Objects 17-99
Configuring LISP Maps 17-102
Configuring M3UA Maps 17-103
M3UA Protocol Conformance 17-103
M3UA Inspection Limitations 17-103
M3UA Policy Maps Add or Edit Match Condition and Action Dialog Boxes 17-105
Configuring Regular Expression Groups 17-108
Add/Edit Regular Expressions 17-108
Metacharacters Used to Build Regular Expressions 17-109
Configuring Settings for Inspection Rules for IOS Devices 17-111

CHAPTER 18 **Managing Firewall Web Filter Rules** 18-1
Understanding Web Filter Rules 18-1
Configuring Web Filter Rules for ASA, PIX, and FWSM Devices 18-2
Web Filter Rules Page (ASA/PIX/FWSM) 18-3
Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes 18-5
Edit Web Filter Type Dialog Box 18-8
Edit Web Filter Options Dialog Box 18-9

Configuring Web Filter Rules for IOS Devices 18-10
Web Filter Rules Page (IOS) 18-12
IOS Web Filter Rule and Applet Scanner Dialog Box 18-13
IOS Web Filter Exclusive Domain Name Dialog Box 18-14
Configuring Settings for Web Filter Servers 18-15
Web Filter Settings Page 18-16
Web Filter Server Configuration Dialog Box 18-19

CHAPTER 19 **Managing Firewall Botnet Traffic Filter Rules** 19-1
Understanding Botnet Traffic Filtering 19-1
Task Flow for Configuring the Botnet Traffic Filter 19-2
Configuring the Dynamic Database 19-4
Adding Entries to the Static Database 19-5
Enabling DNS Snooping 19-6
Enabling Traffic Classification and Actions for the Botnet Traffic Filter 19-6
Botnet Traffic Filter Rules Page 19-9
Dynamic Blacklist Configuration Tab 19-10
Traffic Classification Tab 19-11
BTF Enable Rules Editor 19-12
BTF Drop Rules Editor 19-13
Whitelist/Blacklist Tab 19-14
Device Whitelist or Device Blacklist Dialog Box 19-15

CHAPTER 20 **Working with ScanSafe Web Security** 20-1
Configuring ScanSafe Web Security 20-2
ScanSafe Web Security Page 20-4
Add and Edit Default User Groups Dialog Box 20-6
ScanSafe Web Security Settings Page 20-6

CHAPTER 21 **Managing Zone-based Firewall Rules** 21-1
Understanding the Zone-based Firewall Rules 21-3
The Self Zone 21-5
Using VPNs with Zone-based Firewall Policies 21-6
Zones and VRF-aware Firewalls 21-7
Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules 21-8
Understanding the Relationship Between Services and Protocols in Zone-based Firewall Rules 21-11
General Recommendations for Zone-based Firewall Rules 21-12
Developing and Applying Zone-based Firewall Rules 21-12

Adding Zone-Based Firewall Rules 21-13
Configuring Inspection Maps for Zone-based Firewall Policies 21-16
Configuring Class Maps for Zone-Based Firewall Policies 21-19
Zone-based Firewall IM Application Class Maps: Add or Edit Match Condition Dialog Boxes 21-21
Zone-based Firewall P2P Application Class Maps: Add or Edit Match Condition Dialog Boxes 21-21
H.323 (IOS) Class Maps Add or Edit Match Criterion Dialog Boxes 21-22
HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes 21-22
IMAP and POP3 Class Maps Add or Edit Match Criterion Dialog Boxes 21-25
SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes 21-25
SMTP Class Maps Add or Edit Match Criterion Dialog Boxes 21-27
Sun RPC Class Maps Add or Edit Match Criterion Dialog Boxes 21-29
Local Web Filter Class Add or Edit Match Criterion Dialog Boxes 21-29
N2H2 and Websense Class Add or Edit Match Criterion Dialog Boxes 21-30
Configuring Inspect Parameter Maps 21-31
Configuring Protocol Info Parameter Maps 21-33
Add or Edit DNS Server for Protocol Info Parameters Dialog Box 21-34
Configuring Policy Maps for Zone-Based Firewall Policies 21-34
Add or Edit Match Condition and Action Dialog Boxes for Zone-Based Firewall and Web Filter Policies 21-35
Configuring Content Filtering Maps for Zone-based Firewall Policies 21-36
Configuring Local Web Filter Parameter Maps 21-38
Configuring N2H2 or WebSense Parameter Maps 21-39
Add or Edit External Filter Dialog Box 21-41
Configuring Trend Parameter Maps 21-42
Configuring URL Filter Parameter Maps 21-43
Add or Edit URL Domain Name Dialog Box for URL Filter Parameters 21-45
Configuring URLF Glob Parameter Maps 21-45
Configuring Web Filter Maps 21-47
Changing the Default Drop Behavior 21-48
Configuring Settings for Zone-based Firewall Rules 21-49
Zone Based Firewall Page 21-50
Zone Based Firewall Page - Content Filter Tab 21-52
Zone Dialog Box 21-53
Troubleshooting Zone-based Rules and Configurations 21-54
Zone-based Firewall Rules Page 21-58
Adding and Editing Zone-based Firewall Rules 21-62
Zone-based Firewall Rule: Advanced Options Dialog Box 21-67
Protocol Selector Dialog Box 21-68

Configure Protocol Dialog Box 21-69

CHAPTER 22 **Managing Traffic Zones** 22-1

Why Use Zones? 22-1

ECMP Routing 22-4

Understanding Traffic Zones 22-6

Prerequisites for Traffic Zones 22-7

Guidelines for Traffic Zones 22-8

Configuring Traffic Zones 22-9

CHAPTER 23 **Managing Transparent Firewall Rules** 23-1

Configuring Transparent Firewall Rules 23-1

Transparent Rules Page 23-3

Add and Edit Transparent Firewall Rule Dialog Boxes 23-5

Edit Transparent EtherType Dialog Box 23-7

Edit Transparent Mask Dialog Box 23-7

CHAPTER 24 **Configuring Network Address Translation** 24-1

Understanding Network Address Translation 24-2

Types of Address Translation 24-3

About "Simplified" NAT on ASA 8.3+ Devices 24-4

NAT Policies on Cisco IOS Routers 24-5

NAT Page: Interface Specification 24-6

NAT Page: Static Rules 24-6

NAT Static Rule Dialog Boxes 24-7

NAT Page: Dynamic Rules 24-10

NAT Dynamic Rule Dialog Box 24-11

NAT Page: Timeouts 24-13

NAT Policies on Security Devices 24-15

NAT in Transparent Mode 24-16

Global Options Page 24-16

Translation Options Page 24-17

Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices 24-18

Address Pools 24-18

Translation Rules: PIX, FWSM, and pre-8.3 ASA 24-19

Translation Exemptions (NAT 0 ACL) 24-20

Dynamic Rules Tab 24-22

Policy Dynamic Rules Tab 24-24

Static Rules Tab 24-26
General Tab 24-31
Configuring NAT on ASA 8.3+ Devices 24-33
Translation Rules: ASA 8.3+ 24-34
Per-Session NAT Rules: ASA 9.0(1)+ 24-46

VPN Configuration

CHAPTER 25 Managing Site-to-Site VPNs: The Basics 25-1
Understanding VPN Topologies 25-2
Hub-and-Spoke VPN Topologies 25-2
Point-to-Point VPN Topologies 25-3
Full Mesh VPN Topologies 25-4
Implicitly Supported Topologies 25-5
Understanding IPsec Technologies and Policies 25-5
Understanding Mandatory and Optional Policies for Site-to-Site VPNs 25-6
Overview of Site-to-Site VPN Policies 25-8
Understanding Devices Supported by Each IPsec Technology 25-9
Including Unmanaged or Non-Cisco Devices in a VPN 25-11
Understanding and Configuring VPN Default Policies 25-12
Using Device Overrides to Customize VPN Policies 25-13
Understanding VRF-Aware IPsec 25-14
VRF-Aware IPsec One-Box Solution 25-14
VRF-Aware IPsec Two-Box Solution 25-15
Enabling and Disabling VRF on Catalyst Switches and 7600 Devices 25-17
Accessing Site-to-Site VPN Topologies and Policies 25-17
Site-to-Site VPN Manager Window 25-18
Configuring VPN Topologies in Device View 25-19
Site-To-Site VPN Discovery 25-20
Supported and Unsupported Technologies and Topologies for VPN Discovery 25-20
Prerequisites for VPN Discovery 25-21
VPN Discovery Rules 25-22
Discovering Site-to-Site VPNs 25-24
Defining or Repairing Discovered VPNs with Multiple Spoke Definitions 25-26
Rediscovering Site-to-Site VPNs 25-27
Creating or Editing VPN Topologies 25-28
Defining the Name and IPsec Technology of a VPN Topology 25-30
Selecting Devices for Your VPN Topology 25-32
Defining the Endpoints and Protected Networks 25-34

Configuring VPN Interface Endpoint Settings 25-36
Configuring Dial Backup 25-40
Dial Backup Settings Dialog Box 25-41
Configuring VPNSM or VPN SPA/VSPA Endpoint Settings 25-42
Identifying the Protected Networks for Endpoints 25-46
Configuring a Firewall Services Module (FWSM) Interface with VPNSM or VPNSPA/VSPA 25-47
Configuring VRF Aware IPsec Settings 25-48
Configuring Crypto Map 25-50
Configuring High Availability in Your VPN Topology 25-52
Defining GET VPN Group Encryption 25-54
Add Certificate Filter Dialog Box 25-58
Add New or Edit Security Association Dialog Box 25-58
Defining GET VPN Peers 25-60
Assigning Initial Policies (Defaults) to a New VPN Topology 25-62
Viewing a Summary of a VPN Topology's Configuration 25-63
Creating or Editing Extranet VPNs 25-66
Deleting a VPN Topology 25-71

CHAPTER 26 **Configuring IKE and IPsec Policies** 26-1
Overview of IKE and IPsec Configurations 26-2
Comparing IKE Version 1 and 2 26-4
Understanding IKE 26-5
Deciding Which Encryption Algorithm to Use 26-6
Deciding Which Hash Algorithm to Use 26-6
Deciding Which Diffie-Hellman Modulus Group to Use 26-7
Deciding Which Authentication Method to Use 26-8
Configuring an IKE Proposal 26-9
Configuring IKEv1 Proposal Policy Objects 26-10
Configuring IKEv2 Proposal Policy Objects 26-14
Understanding IPsec Proposals 26-18
Understanding IPsec Proposals for Site-to-Site VPNs 26-19
Understanding Crypto Maps 26-19
Understanding Transform Sets 26-20
Understanding Reverse Route Injection 26-21
Configuring IPsec Proposals in Site-to-Site VPNs 26-22
Selecting the IKE Version for Devices in Site-to-Site VPNs 26-26
Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects 26-27
Configuring VPN Global Settings 26-30

Configuring VPN Global Address Assignment Settings 26-31
Configuring VPN Global ISAKMP/IPsec Settings 26-33
Configuring VPN Global IKEv2 Settings 26-37
Understanding NAT in VPNs 26-41
Configuring VPN Global NAT Settings 26-42
Configuring VPN Global General Settings 26-44
Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs 26-47
Configuring IKEv1 Preshared Key Policies 26-48
Understanding Public Key Infrastructure Policies 26-51
Requirements for Successful PKI Enrollment 26-52
Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs 26-54
Defining Multiple IKEv1 CA Servers for Site-to-Site VPNs 26-55
Configuring Public Key Infrastructure Policies for Remote Access VPNs 26-56
PKI Enrollment Dialog Box 26-58
PKI Enrollment Dialog Box—CA Information Tab 26-60
PKI Enrollment Dialog Box—Enrollment Parameters Tab 26-63
PKI Enrollment Dialog Box—Certificate Subject Name Tab 26-66
PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab 26-67
Configuring IKEv2 Authentication in Site-to-Site VPNs 26-68
IKEv2 Authentication Policy 26-70
IKEv2 Authentication (Override) Dialog Box 26-72

CHAPTER 27 **GRE and DM VPNs** 27-1
Understanding the GRE Modes Page 27-1
GRE and Dynamic GRE VPNs 27-2
Understanding GRE 27-2
Advantages of IPsec Tunneling with GRE 27-3
How Does Security Manager Implement GRE? 27-3
Prerequisites for Successful Configuration of GRE 27-3
Understanding GRE Configuration for Dynamically Addressed Spokes 27-5
Configuring IPsec GRE VPNs 27-5
Configuring GRE Modes for GRE or GRE Dynamic IP VPNs 27-6
Dynamic Multipoint VPNs (DMVPN) 27-9
Understanding DMVPN 27-10
Enabling Spoke-to-Spoke Connections in DMVPN Topologies 27-10
Advantages of DMVPN with GRE 27-11
Configuring DMVPN 27-12
Configuring GRE Modes for DMVPN 27-12
Configuring Large Scale DMVPNs 27-16

Configuring Server Load Balancing in Large Scale DMVPN 27-17
Edit Load Balancing Parameters Dialog Box 27-17

CHAPTER 28 **Easy VPN** 28-1

Understanding Easy VPN 28-1
Easy VPN with Dial Backup 28-2
Easy VPN with High Availability 28-2
Easy VPN with Dynamic Virtual Tunnel Interfaces 28-3
Easy VPN Configuration Modes 28-3
Easy VPN and IKE Extended Authentication (Xauth) 28-4
Overview of Configuring Easy VPN 28-5
Important Notes About Easy VPN Configuration 28-6
Configuring Client Connection Characteristics for Easy VPN 28-7
Configuring Credentials Policy Objects 28-9
Configuring an IPsec Proposal for Easy VPN 28-10
Configuring Dynamic VTI for Easy VPN 28-12
Configuring a Connection Profile Policy for Easy VPN 28-13
Configuring a User Group Policy for Easy VPN 28-14

CHAPTER 29 **Group Encrypted Transport (GET) VPNs** 29-1

Understanding Group Encrypted Transport (GET) VPNs 29-2
Understanding the GET VPN Registration Process 29-4
Choosing the Rekey Transport Mechanism 29-6
Configuring Redundancy Using Cooperative Key Servers 29-7
Configuring Fail-Close to Protect Registration Failures 29-8
Understanding the GET VPN Security Policy and Security Associations 29-10
Understanding Time-Based Anti-Replay 29-11
Configuring GET VPN 29-12
Generating and Synchronizing RSA Keys 29-13
Configuring the IKE Proposal for GET VPN 29-15
Configuring Global Settings for GET VPN 29-16
Configuring GET VPN Key Servers 29-18
Add Key Server, Group Member Dialog Box 29-19
Edit Key Server Dialog Box 29-19
Configuring GET VPN Group Members 29-20
Edit Group Member Dialog Box 29-21
Using Passive Mode to Migrate to GET VPN 29-23
Troubleshooting GET VPN Configurations 29-25

CHAPTER 30 **Managing Remote Access VPNs: The Basics** 30-1

Understanding Remote Access VPNs 30-1
Understanding Remote Access IPSec VPNs 30-2
Understanding Remote Access SSL VPNs 30-2
Remote Access SSL VPN Example 30-3
SSL VPN Access Modes 30-4
Understanding and Managing SSL VPN Support Files 30-5
Prerequisites for Configuring SSL VPNs 30-7
SSL VPN Limitations 30-8
Understanding Devices Supported by Each Remote Access VPN Technology 30-8
Overview of Remote Access VPN Policies 30-9
Discovering Remote Access VPN Policies 30-12
Using the Remote Access VPN Configuration Wizard 30-13
Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices) 30-14
SSL VPN Configuration Wizard—Access Page (ASA) 30-16
SSL VPN Configuration Wizard—Connection Profile Page (ASA) 30-17
Creating User Groups with the Create Group Policy Wizard 30-20
Create Group Policy Wizard—Full Tunnel Page 30-21
Create Group Policy Wizard—Clientless and Thin Client Access Modes Page 30-24
Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices) 30-25
Remote Access VPN Configuration Wizard—IPSec VPN Connection Profile Page (ASA) 30-28
Remote Access VPN Configuration Wizard—IPSec Settings Page (ASA) 30-30
Remote Access VPN Configuration Wizard—Defaults Page 30-31
Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices) 30-32
SSL VPN Configuration Wizard—Gateway and Context Page (IOS) 30-33
SSL VPN Configuration Wizard—Portal Page Customization Page (IOS) 30-35
Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices) 30-36

CHAPTER 31 **Managing Remote Access VPNs on ASA and PIX 7.0+ Devices** 31-1

Overview of Remote Access VPN Policies for ASA and PIX 7.0+ Devices 31-2
Understanding Cluster Load Balancing (ASA) 31-5
Configuring Cluster Load Balance Policies (ASA) 31-5
Configuring Connection Profiles (ASA, PIX 7.0+) 31-7
Connection Profiles Page 31-8
Supported CLIs in Remote Access VPN Multi-Context Mode - Connection Profiles 31-9
General Tab (Connection Profiles) 31-10
AAA Tab (Connection Profiles) 31-13

Secondary AAA Tab (Connection Profiles) **31-17**
IPSec Tab (Connection Profiles) **31-19**
SSL Tab (Connection Profiles) **31-22**
Configuring Group Policies for Remote Access VPNs **31-26**
Understanding Group Policies (ASA) **31-27**
Creating Group Policies (ASA, PIX 7.0+) **31-28**
Understanding SSL VPN Server Verification (ASA) **31-30**
Configuring Trusted Pool Settings (ASA) **31-30**
Using the Trustpool Manager **31-32**
Add/Edit Scripts Dialog Box **31-34**
Working with IPSec VPN Policies **31-36**
Configuring Certificate to Connection Profile Map Policies (ASA) **31-36**
Configuring Certificate to Connection Profile Map Rules (ASA) **31-37**
Map Rule Dialog Box (Upper Table) **31-39**
Map Rule Dialog Box (Lower Table) **31-39**
Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices) **31-40**
IPsec Proposal Editor (ASA, PIX 7.0+ Devices) **31-41**
Working with SSL and IKEv2 IPSec VPN Policies **31-43**
Understanding SSL VPN Access Policies (ASA) **31-44**
SSL VPN Access Policy Page **31-45**
Configuring an Access Policy **31-49**
Configuring Other SSL VPN Settings (ASA) **31-51**
Configuring SSL VPN Performance Settings (ASA) **31-52**
Configuring SSL VPN Content Rewrite Rules (ASA) **31-53**
Configuring SSL VPN Encoding Rules (ASA) **31-55**
Configuring SSL VPN Proxies and Proxy Bypass (ASA) **31-57**
Configuring SSL VPN Browser Plug-ins (ASA) **31-60**
Understanding SSL VPN AnyConnect Client Settings **31-62**
Configuring SSL VPN AnyConnect Client Settings (ASA) **31-64**
Understanding Kerberos Constrained Delegation (KCD) for SSL VPN (ASA) **31-66**
Configuring Kerberos Constrained Delegation (KCD) for SSL VPN (ASA) **31-69**
Configuring AnyConnect Custom Attributes (ASA) **31-70**
Configuring SSL VPN Advanced Settings (ASA) **31-72**
Configuring SSL VPN Server Verification (ASA) **31-73**
Configuring SSL VPN Shared Licenses (ASA 8.2+) **31-74**
Configuring an ASA Device as a Shared License Client **31-76**
Configuring an ASA Device as a Shared License Server **31-77**
Customizing Clientless SSL VPN Portals **31-77**
Configuring ASA Portal Appearance Using SSL VPN Customization Objects **31-78**

Localizing SSL VPN Web Pages for ASA Devices 31-80
Creating Your Own SSL VPN Logon Page for ASA Devices 31-82
Configuring SSL VPN Bookmark Lists for ASA and IOS Devices 31-82
Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks 31-84
Configuring SSL VPN Smart Tunnels for ASA Devices 31-85
Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs 31-88

CHAPTER 32 **Managing Dynamic Access Policies for Remote Access VPNs (ASA 8.0+ Devices) 32-1**
Understanding Dynamic Access Policies 32-1
Configuring Dynamic Access Policies 32-2
Understanding DAP Attributes 32-4
Configuring DAP Attributes 32-7
Configuring Cisco Secure Desktop Policies on ASA Devices 32-9
Dynamic Access Page (ASA) 32-11
Add/Edit Dynamic Access Policy Dialog Box 32-12
Main Tab 32-14
Logical Operations Tab 32-42
Advanced Expressions Tab 32-44
Cisco Secure Desktop Manager Policy Editor Dialog Box 32-46

CHAPTER 33 **Managing Remote Access VPNs on IOS and PIX 6.3 Devices 33-1**
Overview of Remote Access VPN Policies for IOS and PIX 6.3 Devices 33-2
Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices) 33-3
IPsec Proposal Editor (IOS, PIX 6.3 Devices) 33-4
VPNSM/VPN SPA/VSPA Settings Dialog Box 33-6
Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices) 33-7
Configuring High Availability in Remote Access VPNs (IOS) 33-11
Configuring User Group Policies 33-13
Configuring an SSL VPN Policy (IOS) 33-14
SSL VPN Context Editor Dialog Box (IOS) 33-15
General Tab 33-16
Creating Cisco Secure Desktop Configuration Objects 33-18

CHAPTER 34 **Configuring Policy Objects for Remote Access VPNs 34-1**
ASA Group Policies Dialog Box 34-1
Override ASA Group Policy 34-4
Supported CLIs in Remote Access VPN Multi-Context Mode - Group Policy 34-5
ASA Group Policies Client Configuration Settings 34-6

ASA Group Policies Client Firewall Attributes 34-7
ASA Group Policies Hardware Client Attributes 34-9
ASA Group Policies IPSec Settings 34-10
Add or Edit Client Access Rules Dialog Box 34-12
ASA Group Policies SSL VPN Clientless Settings 34-12
Add or Edit VDI Server Dialog Box 34-15
ASA Group Policies SSL VPN Full Client Settings 34-19
ASA Group Policies SSL VPN Settings 34-25
Add or Edit Auto Signon Rules Dialog Box 34-27
ASA Group Policies Browser Proxy Settings 34-29
ASA Group Policies DNS/WINS Settings 34-30
ASA Group Policies Split Tunneling Settings 34-31
ASA Group Policies Connection Settings 34-33
Add or Edit Secure Desktop Configuration Dialog Box 34-35
Add and Edit File Object Dialog Boxes 34-37
File Object — Choose a file Dialog Box 34-39
Add or Edit Port Forwarding List Dialog Boxes 34-40
Add or Edit A Port Forwarding Entry Dialog Box 34-41
Add or Edit Single Sign On Server Dialog Boxes 34-42
Add or Edit Bookmarks Dialog Boxes 34-44
Add or Edit Bookmark Entry Dialog Boxes 34-45
Add and Edit Post Parameter Dialog Boxes 34-48
Add and Edit SSL VPN Customization Dialog Boxes 34-51
SSL VPN Customization Dialog Box—Title Panel 34-53
SSL VPN Customization Dialog Box—Language 34-54
Add and Edit Language Dialog Boxes 34-56
SSL VPN Customization Dialog Box—Logon Form 34-56
SSL VPN Customization Dialog Box—Informational Panel 34-57
SSL VPN Customization Dialog Box—Copyright Panel 34-58
SSL VPN Customization Dialog Box—Full Customization 34-59
SSL VPN Customization Dialog Box—Toolbar 34-59
SSL VPN Customization Dialog Box—Applications 34-60
SSL VPN Customization Dialog Box—Custom Panes 34-61
Add and Edit Column Dialog Boxes 34-61
Add or Edit Custom Pane Dialog Boxes 34-62
SSL VPN Customization Dialog Box—Home Page 34-62
SSL VPN Customization Dialog Box—Logout Page 34-63
Add or Edit SSL VPN Gateway Dialog Box 34-64
Add and Edit Smart Tunnel List Dialog Boxes 34-66

Add and Edit A Smart Tunnel Entry Dialog Boxes 34-67
Add and Edit Smart Tunnel Network Lists Dialog Boxes 34-69
Add and Edit A Smart Tunnel Network List Entry Dialog Box 34-70
Add and Edit Smart Tunnel Auto Signon List Dialog Boxes 34-71
Add and Edit Smart Tunnel Auto Signon Entry Dialog Boxes 34-72
Add or Edit User Group Dialog Box 34-73
User Group Dialog Box—General Settings 34-75
User Group Dialog Box—DNS/WINS Settings 34-77
User Group Dialog Box—Split Tunneling 34-77
User Group Dialog Box—IOS Client Settings 34-78
User Group Dialog Box—IOS Xauth Options 34-80
User Group Dialog Box—IOS Client VPN Software Update 34-81
Add/Edit Client Update Dialog Box 34-81
User Group Dialog Box—Advanced PIX Options 34-82
User Group Dialog Box—Clientless Settings 34-83
User Group Dialog Box—Thin Client Settings 34-84
User Group Dialog Box—SSL VPN Full Tunnel Settings 34-84
User Group Dialog Box—SSL VPN Split Tunneling 34-86
User Group Dialog Box—Browser Proxy Settings 34-87
User Group Dialog Box—SSL VPN Connection Settings 34-88
Add or Edit WINS Server List Dialog Box 34-89
Add or Edit WINS Server Dialog Box 34-90

CHAPTER 35 **Using Map View** 35-1

Understanding Maps and Map View 35-1
Understanding the Map View Main Page 35-2
Map Toolbar 35-4
Using the Navigation Window 35-4
Maps Context Menus 35-5
Managed Device Node Context Menu 35-5
Multiple Selected Nodes Context Menu 35-6
VPN Connection Context Menu 35-6
Layer 3 Link Context Menu 35-7
Map Object Context Menu 35-7
Map Background Context Menu 35-7
Access Permissions for Maps 35-8
Working With Maps 35-8
Creating New or Default Maps 35-9
Opening Maps 35-10

Saving Maps 35-10
Deleting Maps 35-10
Exporting Maps 35-11
Arranging Map Elements 35-11
Panning, Centering, and Zooming Maps 35-11
Selecting Map Elements 35-12
Searching for Map Nodes 35-12
Using Linked Maps 35-13
Setting the Map Background Properties 35-13
Displaying Your Network on the Map 35-14
Understanding Map Elements 35-14
Displaying Managed Devices on the Map 35-16
Showing Containment of Catalyst Switches, Firewalls, and Adaptive Security Appliances 35-16
Using Map Objects To Represent Network Topology 35-17
Add Map Object and Node Properties Dialog Boxes 35-18
Select Policy Object Dialog Box 35-18
Interface Properties Dialog Box 35-19
Creating and Managing Layer 3 Links on the Map 35-19
Select Interfaces and Link Properties Dialog Boxes 35-20
Add Link Dialog Box 35-20
Managing VPNs in Map View 35-20
Displaying Existing VPNs on the Map 35-21
Creating VPN Topologies in Map View 35-21
Editing VPN Policies or Peers From the Map 35-22
Managing Device Policies in Map View 35-22
Performing Basic Policy Management in Map View 35-22
Managing Firewall Policies in Map View 35-23
Managing Firewall Settings in Map View 35-23

IPS Configuration

CHAPTER 36 **Getting Started with IPS Configuration** 36-1
Understanding IPS Network Sensing 36-2
Capturing Network Traffic 36-2
Correctly Deploying the Sensor 36-4
Tuning the IPS 36-4
Overview of IPS Configuration 36-5
Identifying Allowed Hosts 36-7
Configuring SNMP 36-8

General SNMP Configuration Options 36-10
SNMPv3 Users Tab 36-11
Add SNMPv3 User Dialog Box 36-12
SNMP Trap Configuration Tab 36-13
SNMP Trap Communication Dialog Box 36-14
Managing User Accounts and Password Requirements 36-15
Understanding IPS User Roles 36-15
Understanding Managed and Unmanaged IPS Passwords 36-16
Understanding How IPS Passwords are Discovered and Deployed 36-17
Configuring IPS User Accounts 36-18
Add User and Edit User Credentials Dialog Boxes 36-19
Configuring User Password Requirements 36-20
Configuring AAA Access Control for IPS Devices 36-21
Identifying an NTP Server 36-23
Identifying DNS Servers 36-24
Identifying an HTTP Proxy Server 36-24
IPS SSHv2 Known Host Keys 36-25
Add or Edit Known Host RSA Key Dialog Box 36-25
Configuring IPS SSHv1 Fallback Settings 36-26
Configuring the External Product Interface 36-26
External Product Interface Dialog Box 36-27
Posture ACL Dialog Box 36-29
Configuring IPS Logging Policies 36-30
IPS Health Monitor 36-31
Configuring IPS Security Settings 36-32

CHAPTER 37 **Managing IPS Device Interfaces** 37-1
Understanding Interfaces 37-1
Understanding Interface Modes 37-2
Promiscuous Mode 37-2
Inline Interface Mode 37-3
Inline VLAN Pair Mode 37-3
VLAN Group Mode 37-4
Deploying VLAN Groups 37-5
Configuring Interfaces 37-6
Understanding the IPS Interfaces Policy 37-6
Viewing a Summary of IPS Interface Configuration 37-8
Configuring Physical Interfaces 37-9

Modify Physical Interface Map Dialog Box 37-10
Configuring Bypass Mode 37-12
Configuring CDP Mode 37-12
Configuring Inline Interface Pairs 37-13
Configuring Inline VLAN Pairs 37-14
Configuring VLAN Groups 37-15

CHAPTER 38 **Configuring Virtual Sensors** 38-1
Understanding the Virtual Sensor 38-1
Advantages and Restrictions of Virtualization 38-3
Inline TCP Session Tracking Mode 38-3
Understanding Normalizer Mode 38-4
Assigning Interfaces to Virtual Sensors 38-4
Identifying the Virtual Sensors for a Device 38-5
Defining A Virtual Sensor 38-5
Virtual Sensor Dialog Box 38-7
Editing Policies for a Virtual Sensor 38-9
Deleting A Virtual Sensor 38-10

CHAPTER 39 **Defining IPS Signatures** 39-1
Understanding Signatures 39-1
Obtaining Detailed Information About a Signature 39-2
Understanding Signature Inheritance 39-3
IPS Signature Purge 39-3
Configuring Signatures 39-4
Signatures Page 39-4
Apply Signature Threat Profiles 39-9
Signature Shortcut Menu 39-10
Edit, Add, Replace Action Dialog Boxes 39-12
Edit Fidelity Dialog Box 39-13
Viewing Signature Update Levels 39-13
Enabling and Disabling Signatures 39-14
Editing Signatures 39-14
Edit Signature or Add Custom Signature Dialog Boxes 39-15
Adding Custom Signatures 39-19
Engine Options 39-20
Cloning Signatures 39-21
Regular Expressions in Custom Signatures 39-22
Editing Signature Parameters (Tuning Signatures) 39-23

Edit Signature Parameters Dialog Box 39-24
Editing the Component List for Meta Engine Signatures 39-29
Obsoletes Dialog Box 39-30
Configuring Signature Settings 39-30

CHAPTER 40 **Configuring Event Action Rules** 40-1
Understanding the IPS Event Action Process 40-1
Understanding IPS Event Actions 40-2
Configuring Event Action Filters 40-4
Tips for Managing Event Action Filter Rules 40-6
Event Action Filters Page 40-7
Filter Item Dialog Box 40-9
Configuring Event Action Overrides 40-13
Add or Edit Event Action Rule Dialog Box 40-14
Configuring Risk Rating Policy Objects 40-15
Add or Edit Risk Rating Dialog Box 40-16
Configuring IPS Event Action Network Information 40-17
Configuring Target Value Ratings 40-17
Target Value Rating Dialog Box 40-19
Understanding Passive OS Fingerprinting 40-19
Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only) 40-21
OS Map Dialog Box 40-22
Configuring Settings for Event Actions 40-23

CHAPTER 41 **Managing IPS Anomaly Detection** 41-1
Understanding Anomaly Detection 41-1
Worm Viruses 41-2
Anomaly Detection Modes 41-2
Anomaly Detection Zones 41-3
Knowing When to Turn Off Anomaly Detection 41-4
Configuring Anomaly Detection Signatures 41-4
Configuring Anomaly Detection 41-6
Configuring Anomaly Detection Learning Accept Mode 41-8
Understanding Anomaly Detection Thresholds and Histograms 41-9
Configuring Anomaly Detection Thresholds and Histograms 41-11
Dest Port or Protocol Map Dialog Box 41-12
Histogram Dialog Box 41-13

CHAPTER 42 **Configuring Global Correlation** 42-1

Understanding Global Correlation 42-1
Understanding Reputation 42-2
Understanding Network Participation 42-3
Global Correlation Requirements and Limitations 42-4
Configuring Global Correlation Inspection and Reputation 42-5
Configuring Network Participation 42-7

CHAPTER 43 **Configuring Attack Response Controller for Blocking and Rate Limiting** 43-1

Understanding IPS Blocking 43-1
Strategies for Applying Blocks 43-3
Understanding Rate Limiting 43-4
Understanding Router and Switch Blocking Devices 43-4
Understanding the Master Blocking Sensor 43-6
Configuring IPS Blocking and Rate Limiting 43-7
Blocking Page 43-8
General Tab, IPS Blocking Policy 43-10
User Profile Dialog Box 43-12
Master Blocking Sensor Dialog Box 43-13
Router, Firewall, Cat6K Device Dialog Box 43-14
Router Block Interface Dialog Box 43-15
Cat6k Block VLAN Dialog Box 43-16
Never Block Host or Network Dialog Boxes 43-17

CHAPTER 44 **Managing IPS Sensors** 44-1

Managing IPS Licenses 44-1
Updating IPS License Files 44-1
Redeploying IPS License Files 44-2
Automating IPS License File Updates 44-3
Managing IPS Updates 44-4
Configuring the IPS Update Server 44-4
Checking for IPS Updates and Downloading Them 44-5
Automating IPS Updates 44-6
Manually Applying IPS Updates 44-7
Managing IPS Certificates 44-10
Rebooting IPS Sensors 44-12

CHAPTER 45 **Configuring IOS IPS Routers** 45-1

Understanding Cisco IOS IPS 45-1
Understanding IPS Subsystems and Support of IOS IPS Revisions 45-2
Cisco IOS IPS Signature Scanning with Lightweight Signatures 45-2
Router Configuration Files and Signature Event Action Processor (SEAP) 45-3
Cisco IOS IPS Limitations and Restrictions 45-3
Overview of Cisco IOS IPS Configuration 45-4
Initial Preparation of a Cisco IOS IPS Router 45-5
Selecting a Signature Category for Cisco IOS IPS 45-6
Configuring General Settings for Cisco IOS IPS 45-7
Configuring IOS IPS Interface Rules 45-9
IPS Rule Dialog Box 45-10
Pair Dialog Box 45-11

PIX/ASA/FWSM Device Configuration

CHAPTER 46 **Managing Firewall Devices** 46-1

Firewall Device Types 46-1
Default Firewall Configurations 46-2
Configuring Firewall Device Interfaces 46-3
Understanding Device Interfaces 46-3
Interfaces in Routed and Transparent Modes 46-5
Interfaces in Single and Multiple Contexts 46-5
About Asymmetric Routing Groups 46-6
Understanding ASA 5505 Ports and Interfaces 46-6
Configuring Subinterfaces (PIX/ASA) 46-7
Configuring Redundant Interfaces 46-8
Configuring EtherChannels 46-9
Configuring VNI Interfaces 46-15
Configuring Tunnel Interface 46-22
Establishing Regular IPSec VPN Tunnel 46-24
Configuring IPSec Policy for Tunnel Interface 46-24
Managing Device Interfaces, Hardware Ports, and Bridge Groups 46-26
Add/Edit Interface Dialog Box (PIX 6.3) 46-28
Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM) 46-31
Add/Edit Interface Dialog Box: Cisco Firepower 9000 (General and Advanced tabs) 46-40
Configuring Hardware Ports on an ASA 5505 46-61
Add/Edit Bridge Group Dialog Box 46-62
Advanced Interface Settings (PIX/ASA/FWSM) 46-68

Enabling Traffic between Interfaces with the Same Security Level 46-70
Managing the PPPoE Users List 46-71
Managing VPDN Groups 46-72
VXLAN 46-73
Configuring VXLAN Policy 46-73

CHAPTER 47 **Configuring Bridging Policies on Firewall Devices** 47-1
About Bridging on Firewall Devices 47-1
Bridging Support for FWSM 3.1 47-3
ARP Table Page 47-3
Add/Edit ARP Configuration Dialog Box 47-5
ARP Inspection Page 47-5
Add/Edit ARP Inspection Dialog Box 47-6
Managing the IPv6 Neighbor Cache 47-7
MAC Address Table Page 47-8
Add/Edit MAC Table Entry Dialog Box 47-8
MAC Learning Page 47-9
Add/Edit MAC Learning Dialog Box 47-9
Management IP Page 47-10
Management IPv6 Page (ASA 5505) 47-11

CHAPTER 48 **Configuring Device Administration Policies on Firewall Devices** 48-1
About AAA on Security Devices 48-1
Preparing for AAA 48-2
Local Database 48-3
AAA for Device Administration 48-4
AAA for Network Access 48-4
AAA for VPN Access 48-4
Configuring AAA - Authentication Tab 48-5
Authorization Tab 48-7
Accounting Tab 48-8
Configuring Banners 48-9
Configuring Boot Image/Configuration Settings 48-10
Images Dialog Box 48-12
Configuring CLI Prompt 48-12
Setting the Device Clock 48-14
Enabling/Disabling FIPS 48-15

Configuring Umbrella Global Policy 48-16
Configuring Device Credentials 48-17
Managing Mount Points 48-18
Add/Edit Mount Point Configuration Dialog Box 48-19
IP Client 48-20
Add/Edit IP Client Dialog Box 48-20
App Agent 48-21

CHAPTER 49 **Configuring Device Access Settings on Firewall Devices** 49-1
Configuring Console Timeout 49-1
HTTP Page 49-2
HTTP Configuration Dialog Box 49-3
Configuring ICMP 49-4
Add and Edit ICMP Dialog Boxes 49-5
Configuring Management Access 49-6
Configuring Secure Shell Access 49-7
Add and Edit SSH Host Dialog Boxes 49-8
Configuring SSL - Basic and Advanced tabs 49-8
Reference Identities 49-13
Add/Edit Reference Identity Dialog Box 49-13
Configuring SNMP 49-14
SNMP Terminology 49-15
SNMP Version 3 49-15
SNMP Page 49-17
SNMP Trap Configuration Dialog Box 49-19
Add/Edit SNMP Host Access Entry Dialog Box 49-22
Add/Edit SNMP Host Group Entry Dialog Box 49-23
Add/Edit SNMP Group Entry Dialog Box 49-24
Add/Edit SNMP User Entry Dialog Box 49-25
Add/Edit SNMP User List Entry Dialog Box 49-27
Telnet Page 49-29
Telnet Configuration Dialog Box 49-29

CHAPTER 50 **Configuring Failover** 50-1
Understanding Failover 50-1
Active/Active Failover 50-3
Stateful Failover 50-4
Basic Failover Configuration 50-5

Adding A Security Context to Failover Group 2 50-8
Additional Steps for an Active/Standby Failover Configuration 50-9
Exporting the Certificate to a File or PKCS12 data 50-9
Importing the Certificate onto the Standby Device 50-9
Failover Policies 50-10
Failover Page (PIX 6.3) 50-10
Edit Failover Interface Configuration Dialog Box (PIX 6.3) 50-12
Failover Page (FWSM) 50-13
Advanced Settings Dialog Box 50-16
Failover Page (ASA/PIX 7.0+) 50-17
Settings Dialog Box 50-21
Failover Page (Security Context) 50-26
Bootstrap Configuration for LAN Failover Dialog Box 50-26

CHAPTER 51 **Configuring Hostname, Resources, User Accounts, and SLAs** 51-1
Hostname Page 51-1
Resource Management on Multi-context FWSMs 51-2
Resources Page 51-3
Add and Edit Resource Dialog Boxes 51-4
Configuring User Accounts 51-7
Add/Edit User Account Dialog Boxes 51-7
Monitoring Service Level Agreements (SLAs) To Maintain Connectivity 51-8
Creating Service Level Agreements 51-9
Configuring SLA Monitor Objects 51-10

CHAPTER 52 **Configuring Server Access Settings on Firewall Devices** 52-1
AUS Page 52-1
Add and Edit Auto Update Server Dialog Boxes 52-3
DHCP Relay Page 52-5
Add and Edit DHCP Relay Agent Configuration Dialog Boxes 52-6
Add and Edit DHCP Relay Server Configuration Dialog Boxes 52-7
DHCP Relay IPv6 Page 52-7
Add and Edit DHCP Relay IPv6 Agent Configuration Dialog Boxes 52-9
Add and Edit DHCP Relay IPv6 Server Configuration Dialog Boxes 52-9
Configuring DHCP Servers 52-10
DHCP Server Page 52-10
Add and Edit DHCP Server Interface Configuration Dialog Boxes 52-12
Add/Edit DHCP Server Advanced Configuration Dialog Box 52-13

DNS Page 52-14
Add DNS Server Group Dialog Box 52-16
Add DNS Server Dialog Box 52-17
Configuring DDNS 52-18
Add/Edit DDNS Interface Rule Dialog Box 52-19
DDNS Update Methods Dialog Box 52-19
NTP Page 52-21
NTP Server Configuration Dialog Box 52-21
SMTP Server Page 52-22
TFTP Server Page 52-23

CHAPTER 53 **Configuring FXOS Server Access Settings on Firepower 2100 Series Devices** 53-1
HTTPS Page 53-1
Add and Edit HTTPS Dialog Boxes 53-2
SSH Page 53-2
Add and Edit SSH Dialog Boxes 53-3
SNMP Page 53-4
Add and Edit SNMP Dialog Boxes 53-4

CHAPTER 54 **Configuring Logging Policies on Firewall Devices** 54-1
NetFlow Page 54-1
Add and Edit Collector Dialog Boxes (NetFlow) 54-2
Embedded Event Manager 54-3
Add and Edit Applet Dialog Boxes 54-5
Add and Edit Syslog Configuration Dialog Boxes 54-7
Add and Edit Action Configuration Dialog Boxes 54-7
E-Mail Setup Page 54-8
Add/Edit Email Recipient Dialog Box 54-8
Event Lists Page 54-9
Message Classes and Associated Message ID Numbers 54-9
Add/Edit Event List Dialog Box 54-10
Add/Edit Syslog Class Dialog Box 54-11
Add/Edit Syslog Message ID Filter Dialog Box 54-11
Logging Filters Page 54-12
Edit Logging Filters Dialog Box 54-13
Configuring Logging Setup 54-14
Logging Setup Page 54-15
Configuring Rate Limit Levels 54-17

Rate Limit Page 54-18
Add/Edit Rate Limit for Syslog Logging Levels Dialog Box 54-18
Add/Edit Rate Limited Syslog Message Dialog Box 54-19
Configuring Syslog Server Setup 54-20
Syslog Relay Configuration 54-21
Server Setup Page 54-21
Logging Levels 54-24
Add/Edit Syslog Message Dialog Box 54-25
Defining Syslog Servers 54-26
Syslog Servers Page 54-27
Add/Edit Syslog Server Dialog Box 54-28

CHAPTER 55 **Configuring Multicast Policies on Firewall Devices** 55-1
Enabling PIM and IGMP 55-1
Configuring IGMP 55-2
IGMP Page - Protocol Tab 55-3
Configure IGMP Parameters Dialog Box 55-4
IGMP Page - Access Group Tab 55-5
Configure IGMP Access Group Parameters Dialog Box 55-5
IGMP Page - Static Group Tab 55-6
Configure IGMP Static Group Parameters Dialog Box 55-6
IGMP Page - Join Group Tab 55-7
Configure IGMP Join Group Parameters Dialog Box 55-7
Configuring Multicast Routes 55-8
Add/Edit MRoute Configuration Dialog Box 55-8
Configuring Multicast Boundary Filters 55-9
Add/Edit MBoundary Configuration Dialog Box 55-9
Add/Edit MBoundary Interface Configuration Dialog Box 55-10
Configuring PIM 55-11
PIM Page - Protocol Tab 55-11
Add/Edit PIM Protocol Dialog Box 55-12
PIM Page - Neighbor Filter Tab 55-12
Add/Edit PIM Neighbor Filter Dialog Box 55-13
PIM Page - Bidirectional Neighbor Filter Tab 55-13
Add/Edit PIM Bidirectional Neighbor Filter Dialog Box 55-14
PIM Page - Rendezvous Points Tab 55-15
Add/Edit Rendezvous Point Dialog Box 55-16
PIM Page - Route Tree Tab 55-17
PIM Page - Request Filter Tab 55-18

Add/Edit Multicast Group Rules Dialog Box 55-19
PIM Page - Bootstrap Router Tab 55-20
Add/Edit Bootstrap Router Dialog Box 55-21

CHAPTER 56 **Configuring Routing Policies on Firewall Devices** 56-1
Configuring No Proxy ARP 56-1
Configuring BGP 56-2
About BGP 56-3
General Tab 56-5
IPv4 Family Tab 56-6
IPv4 Family - General Tab 56-7
Add/Edit Aggregate Address Dialog Box 56-9
Add/Edit Filter Dialog Box 56-10
Add/Edit Neighbor Dialog Box 56-11
Add/Edit Network Dialog Box 56-17
Add/Edit Redistribution Dialog Box 56-18
Add/Edit Route Injection Dialog Box 56-19
IPv6 Family Tab 56-20
IPv6 Family - General Tab 56-21
Add/Edit Aggregate Address Dialog Box 56-22
Add/Edit Neighbor Dialog Box 56-24
Add/Edit Network Dialog Box 56-29
Add/Edit Redistribution Dialog Box 56-30
Add/Edit Route Injection Dialog Box 56-31
Configuring EIGRP 56-32
About EIGRP 56-33
EIGRP Advanced Dialog Box 56-34
Setup Tab 56-36
Filter Rules Tab 56-39
Add/Edit EIGRP Filter Rule Dialog Box 56-40
Neighbors Tab 56-41
Add/Edit EIGRP Neighbor Dialog Box 56-42
Redistribution Tab 56-42
Add/Edit EIGRP Redistribution Dialog Box 56-44
Summary Address Tab 56-45
Add/Edit EIGRP Summary Address Dialog Box 56-46
Interfaces Tab 56-47
Add/Edit EIGRP Interface Dialog Box 56-48
Configuring ISIS 56-49

About ISIS 56-49
General Tab 56-49
IPv4 Family Tab 56-51
IPv4 Family Tab—General Tab 56-51
IPv4 Family Tab—SPF Tab 56-52
IPv4 Family Tab—Redistribution Tab 56-53
IPv6 Family Tab 56-55
IPv6 Family Tab—General Tab 56-55
IPv6 Family Tab—SPF Tab 56-55
IPv6 Family Tab—Redistribution Tab 56-57
IPv6 Family Tab—Summary Prefix 56-57
Authentication Tab 56-58
Link State Packet Tab 56-59
Summary Address Tab 56-60
Network Entity Title Tab 56-61
Interface Tab 56-61
Interface Tab—General Tab 56-62
Interface Tab—Authentication Tab 56-63
Interface Tab—Hello Padding Tab 56-64
Interface Tab—LSP Settings Tab 56-65
Interface Tab—Metrics Tab 56-65
Passive Interfaces Tab 56-66
Configuring BFD Routing 56-66
About BFD 56-66
BFD Asynchronous Mode and Echo Function 56-67
BFD Session Establishment 56-67
BFD Timer Negotiation 56-69
BFD Failure Detection 56-69
BFD Deployment Scenarios 56-70
Create BFD Template 56-70
Add/ Edit BFD Map Dialog Box 56-72
Add/ Edit BFD Interface Dialog Box 56-73
Configuring OSPF 56-75
About OSPF 56-75
General Tab 56-76
OSPF Advanced Dialog Box 56-77
Area Tab 56-81
Add/Edit Area/Area Networks Dialog Box 56-82
Range Tab 56-84
Add/Edit Area Range Network Dialog Box 56-84

Neighbors Tab 56-85
Add/Edit Static Neighbor Dialog Box 56-85
Redistribution Tab 56-86
Redistribution Dialog Box 56-87
Virtual Link Tab 56-89
Add/Edit OSPF Virtual Link Configuration Dialog Box 56-90
Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box 56-91
Filtering Tab 56-92
Add/Edit Filtering Dialog Box 56-93
Filter Rule Tab 56-94
Add/Edit Filter Rule Dialog Box 56-94
Summary Address Tab 56-95
Add/Edit Summary Address Dialog Box 56-96
Interface Tab 56-96
Add/Edit Interface Dialog Box 56-98
Configuring OSPFv3 56-100
About OSPFv3 56-101
Process Tab 56-103
OSPFv3 Advanced Properties Dialog Box 56-104
Area Tab (OSPFv3) 56-108
Add/Edit Redistribution Dialog Box (OSPFv3) 56-112
Add/Edit Summary Prefix Dialog Box (OSPFv3) 56-113
OSPFv3 Interface Tab 56-114
Add/Edit Interface Dialog Box (OSPFv3) 56-114
Add/Edit Neighbor Dialog Box (OSPFv3) 56-118
Configuring RIP 56-119
RIP Page for PIX/ASA 6.3–7.1 and FWSM 56-120
Add/Edit RIP Configuration (PIX/ASA 6.3–7.1 and FWSM) Dialog Boxes 56-121
RIP Page for PIX/ASA 7.2 and Later 56-122
RIP - Setup Tab 56-123
RIP - Redistribution Tab 56-125
RIP - Filtering Tab 56-126
RIP - Interface Tab 56-127
Configuring Static Routes 56-128
Add/Edit Static Route Dialog Box 56-130
Add/Edit IPv6 Static Route Dialog Box 56-131
Configuring Policy Objects for ASA Routing Policies 56-132
Understanding Route Map Objects 56-132
Add or Edit Route Map Object Dialog Boxes 56-136

Add or Edit Policy List Object Dialog Box 56-143
Add or Edit Prefix List Object Dialog Box 56-146
Add or Edit Prefix List Entry Dialog Box 56-148
Add or Edit Prefix List IPv6 Object Dialog Box 56-148
Add or Edit IPv6 Prefix List Entry Dialog Box 56-150
Add or Edit As Path Object Dialog Boxes 56-151
Add or Edit As Path Entry Dialog Box 56-152
Add or Edit Community List Object Dialog Box 56-153
Add or Edit Community List Entry Dialog Box 56-154
Add or Edit Community List Entry Dialog Box 56-155

CHAPTER 57 **Configuring Security Policies on Firewall Devices** 57-1
General Page 57-1
Configuring Floodguard, Anti-Spoofing and Fragment Settings 57-2
Add/Edit General Security Configuration Dialog Box 57-3
Configuring Timeouts 57-4

CHAPTER 58 **Configuring Service Policy Rules on Firewall Devices** 58-1
About Service Policy Rules 58-1
About TCP State Bypass 58-3
Priority Queues Page 58-4
Priority Queue Configuration Dialog Box 58-4
Service Policy Rules Page 58-5
Insert/Edit Service Policy (MPC) Rule Wizard 58-6
Step 1. Configure a Service Policy 58-6
Step 2. Configure the traffic class 58-7
Step 3. Configure the MPC actions 58-8
About IPS Modules on ASA Devices 58-15
About the ASA CX 58-17
ASA CX Auth Proxy Configuration 58-17
Configuring Traffic Flow Objects 58-18
Default Inspection Traffic 58-20
Configuring TCP Maps 58-22
Add and Edit TCP Option Range Dialog Boxes 58-25

CHAPTER 59 **Configuring Security Contexts on Firewall Devices** 59-1
Enabling and Disabling Multiple-Context Mode 59-1
Checklist for Configuring Multiple Security Contexts 59-3

Managing Security Contexts 59-7
Add/Edit Security Context Dialog Box (FWSM) 59-8
Add/Edit Security Context Dialog Box (PIX/ASA) 59-9
Allocate Interfaces Dialog Box (PIX/ASA only) 59-12

CHAPTER 60 **User Preferences** 60-1
Configuring Deployment Preferences on Firewall Devices 60-1
Configuring Transactional Commit Preferences on Firewall Devices 60-2

Router and Switch Device Configuration

CHAPTER 61 **Managing Routers** 61-1
Configuring Routers Running IOS Software Releases 12.1 and 12.2 61-3
Discovering Router Policies 61-3

CHAPTER 62 **Configuring Router Interfaces** 62-1
Basic Interface Settings on Cisco IOS Routers 62-1
Available Interface Types 62-2
Defining Basic Router Interface Settings 62-4
Deleting a Cisco IOS Router Interface 62-6
Router Interfaces Page 62-7
Create Router Interface Dialog Box 62-8
Interface Auto Name Generator Dialog Box 62-12
Advanced Interface Settings on Cisco IOS Routers 62-13
Understanding Helper Addresses 62-14
Advanced Interface Settings Page 62-16
Advanced Interface Settings Dialog Box 62-16
IPS Module Interface Settings on Cisco IOS Routers 62-22
IPS Module Interface Settings Page 62-23
IPS Monitoring Information Dialog Box 62-24
CEF Interface Settings on Cisco IOS Routers 62-25
CEF Interface Settings Page 62-26
CEF Interface Settings Dialog Box 62-27
Dialer Interfaces on Cisco IOS Routers 62-28
Defining Dialer Profiles 62-28
Defining BRI Interface Properties 62-30
Dialer Policy Page 62-31
Dialer Profile Dialog Box 62-32

Dialer Physical Interface Dialog Box 62-33

ADSL on Cisco IOS Routers 62-34

Supported ADSL Operating Modes 62-35

Defining ADSL Settings 62-36

ADSL Policy Page 62-37

ADSL Settings Dialog Box 62-38

SHDSL on Cisco IOS Routers 62-41

Defining SHDSL Controllers 62-41

SHDSL Policy Page 62-42

SHDSL Controller Dialog Box 62-43

Controller Auto Name Generator Dialog Box 62-46

PVCs on Cisco IOS Routers 62-47

Understanding Virtual Paths and Virtual Channels 62-47

Understanding ATM Service Classes 62-48

Understanding ATM Management Protocols 62-49

Understanding ILMI 62-50

Understanding OAM 62-51

Defining ATM PVCs 62-51

Defining OAM Management on ATM PVCs 62-54

PVC Policy Page 62-55

PVC Dialog Box 62-56

PVC Dialog Box—Settings Tab 62-58

PVC Dialog Box—QoS Tab 62-61

PVC Dialog Box—Protocol Tab 62-64

Define Mapping Dialog Box 62-65

PVC Advanced Settings Dialog Box 62-66

PVC Advanced Settings Dialog Box—OAM Tab 62-67

PVC Advanced Settings Dialog Box—OAM-PVC Tab 62-69

PPP on Cisco IOS Routers 62-71

Understanding Multilink PPP (MLP) 62-71

Defining PPP Connections 62-72

Defining Multilink PPP Bundles 62-75

PPP/MLP Policy Page 62-76

PPP Dialog Box 62-77

PPP Dialog Box—PPP Tab 62-78

PPP Dialog Box—MLP Tab 62-80

CHAPTER 63 **Router Device Administration** 63-1

AAA on Cisco IOS Routers 63-2

Supported Authorization Types 63-2
Supported Accounting Types 63-3
Understanding Method Lists 63-3
Defining AAA Services 63-4
AAA Policy Page 63-6
AAA Page—Authentication Tab 63-6
AAA Page—Authorization Tab 63-8
Command Authorization Dialog Box 63-10
AAA Page—Accounting Tab 63-10
Command Accounting Dialog Box 63-13
User Accounts and Device Credentials on Cisco IOS Routers 63-14
Defining Accounts and Credential Policies 63-14
Accounts and Credential s Policy Page 63-16
User Account Dialog Box 63-17
Bridging on Cisco IOS Routers 63-18
Bridge-Group Virtual Interfaces 63-19
Defining Bridge Groups 63-20
Bridging Policy Page 63-21
Bridge Group Dialog Box 63-21
Time Zone Settings on Cisco IOS Routers 63-22
Defining Time Zone and DST Settings 63-22
Clock Policy Page 63-23
CPU Utilization Settings on Cisco IOS Routers 63-25
Defining CPU Utilization Settings 63-25
CPU Policy Page 63-26
HTTP and HTTPS on Cisco IOS Routers 63-28
Defining HTTP Policies 63-29
HTTP Policy Page 63-31
HTTP Page—Setup Tab 63-31
HTTP Page—AAA Tab 63-32
Command Authorization Override Dialog Box 63-34
Line Access on Cisco IOS Routers 63-35
Defining Console Port Setup Parameters 63-35
Defining Console Port AAA Settings 63-37
Defining VTY Line Setup Parameters 63-38
Defining VTY Line AAA Settings 63-40
Console Policy Page 63-42
Console Page—Setup Tab 63-42

Console Page—Authentication Tab 63-44
Console Page—Authorization Tab 63-45
Console Page—Accounting Tab 63-47
VTY Policy Page 63-50
VTY Line Dialog Box 63-51
VTY Line Dialog Box—Setup Tab 63-52
VTY Line Dialog Box—Authentication Tab 63-55
VTY Line Dialog Box—Authorization Tab 63-56
VTY Line Dialog Box—Accounting Tab 63-57
Command Authorization Dialog Box—Line Access 63-60
Command Accounting Dialog Box—Line Access 63-61
Optional SSH Settings on Cisco IOS Routers 63-63
Defining Optional SSH Settings 63-63
Secure Shell Policy Page 63-64
SNMP on Cisco IOS Routers 63-66
Defining SNMP Agent Properties 63-67
Enabling SNMP Traps 63-68
SNMP Policy Page 63-69
Permission Dialog Box 63-70
Trap Receiver Dialog Box 63-71
SNMP Traps Dialog Box 63-72
DNS on Cisco IOS Routers 63-74
Defining DNS Policies 63-75
DNS Policy Page 63-76
IP Host Dialog Box 63-76
Hostnames and Domain Names on Cisco IOS Routers 63-77
Defining Hostname Policies 63-77
Hostname Policy Page 63-78
Memory Settings on Cisco IOS Routers 63-78
Defining Router Memory Settings 63-78
Memory Policy Page 63-79
Secure Device Provisioning on Cisco IOS Routers 63-81
Contents of Bootstrap Configuration 63-82
Secure Device Provisioning Workflow 63-82
Defining Secure Device Provisioning Policies 63-83
Configuring a AAA Server Group for Administrative Introducers 63-84
Secure Device Provisioning Policy Page 63-85
DHCP on Cisco IOS Routers 63-87

Understanding DHCP Database Agents 63-88
Understanding DHCP Relay Agents 63-88
Understanding DHCP Option 82 63-89
Understanding Secured ARP 63-89
Defining DHCP Policies 63-90
Defining DHCP Address Pools 63-91
DHCP Policy Page 63-92
DHCP Database Dialog Box 63-94
IP Pool Dialog Box 63-94
NTP on Cisco IOS Routers 63-96
Defining NTP Servers 63-97
NTP Policy Page 63-98
NTP Server Dialog Box 63-99

CHAPTER 64 **Configuring Identity Policies** 64-1
802.1x on Cisco IOS Routers 64-1
Understanding 802.1x Device Roles 64-2
802.1x Interface Authorization States 64-2
Topologies Supported by 802.1x 64-3
Defining 802.1x Policies 64-4
802.1x Policy Page 64-5
Network Admission Control on Cisco IOS Routers 64-8
Router Platforms Supporting NAC 64-8
Understanding NAC Components 64-9
Understanding NAC System Flow 64-9
Defining NAC Setup Parameters 64-10
Defining NAC Interface Parameters 64-11
Defining NAC Identity Parameters 64-13
Network Admission Control Policy Page 64-14
Network Admission Control Page—Setup Tab 64-14
Network Admission Control Page—Interfaces Tab 64-16
NAC Interface Configuration Dialog Box 64-17
Network Admission Control Page—Identities Tab 64-18
NAC Identity Profile Dialog Box 64-19
NAC Identity Action Dialog Box 64-19

CHAPTER 65 **Configuring Logging Policies** 65-1
Logging on Cisco IOS Routers 65-1
Defining Syslog Logging Setup Parameters 65-1

Defining Syslog Servers 65-3
Understanding Log Message Severity Levels 65-4
NetFlow on Cisco IOS Routers 65-5
Defining NetFlow Parameters 65-6
Syslog Logging Setup Policy Page 65-7
Syslog Servers Policy Page 65-10
Syslog Server Dialog Box 65-11
NetFlow Policy Page 65-12
Adding and Editing NetFlow Interface Settings 65-15

CHAPTER 66 **Configuring Quality of Service** 66-1
Quality of Service on Cisco IOS Routers 66-1
Quality of Service and CEF 66-2
Understanding Matching Parameters 66-2
Understanding Marking Parameters 66-3
Understanding Queuing Parameters 66-4
Tail Drop vs. WRED 66-4
Low-Latency Queuing 66-5
Default Class Queuing 66-6
Understanding Policing and Shaping Parameters 66-6
Understanding the Token-Bucket Mechanism 66-8
Understanding Control Plane Policing 66-9
Defining QoS Policies 66-10
Defining QoS on Interfaces 66-10
Defining QoS on the Control Plane 66-12
Defining QoS Class Matching Parameters 66-13
Defining QoS Class Marking Parameters 66-15
Defining QoS Class Queuing Parameters 66-16
Defining QoS Class Policing Parameters 66-17
Defining QoS Class Shaping Parameters 66-18
Quality of Service Policy Page 66-19
QoS Policy Dialog Box 66-21
QoS Class Dialog Box 66-23
QoS Class Dialog Box—Matching Tab 66-24
Edit ACLs Dialog Box—QoS Classes 66-25
QoS Class Dialog Box—Marking Tab 66-26
QoS Class Dialog Box—Queuing and Congestion Avoidance Tab 66-27
QoS Class Dialog Box—Policing Tab 66-29
QoS Class Dialog Box—Shaping Tab 66-31

CHAPTER 67 **Configuring Routing Policies** 67-1

BGP Routing on Cisco IOS Routers 67-1
Defining BGP Routes 67-2
Redistributing Routes into BGP 67-3
BGP Routing Policy Page 67-4
BGP Page—Setup Tab 67-5
Neighbors Dialog Box 67-6
BGP Page—Redistribution Tab 67-7
BGP Redistribution Mapping Dialog Box 67-7
EIGRP Routing on Cisco IOS Routers 67-8
Defining EIGRP Routes 67-9
Defining EIGRP Interface Properties 67-10
Redistributing Routes into EIGRP 67-12
EIGRP Routing Policy Page 67-13
EIGRP Page—Setup Tab 67-13
EIGRP Setup Dialog Box 67-14
EIGRP Page—Interfaces Tab 67-15
EIGRP Interface Dialog Box 67-16
EIGRP Page—Redistribution Tab 67-17
EIGRP Redistribution Mapping Dialog Box 67-18
OSPF Routing on Cisco IOS Routers 67-19
Defining OSPF Process Settings 67-20
Defining OSPF Area Settings 67-21
Redistributing Routes into OSPF 67-22
Defining OSPF Redistribution Mappings 67-22
Defining OSPF Maximum Prefix Values 67-24
Defining OSPF Interface Settings 67-25
Understanding Interface Cost 67-26
Understanding Interface Priority 67-26
Disabling MTU Mismatch Detection 67-27
Blocking LSA Flooding 67-28
Understanding OSPF Timer Settings 67-28
Understanding the OSPF Network Type 67-29
Understanding OSPF Interface Authentication 67-29
OSPF Interface Policy Page 67-30
OSPF Interface Dialog Box 67-31
OSPF Process Policy Page 67-34
OSPF Process Page—Setup Tab 67-35
OSPF Setup Dialog Box 67-35

Edit Interfaces Dialog Box—OSPF Passive Interfaces 67-36
OSPF Process Page—Area Tab 67-36
OSPF Area Dialog Box 67-37
OSPF Process Page—Redistribution Tab 67-38
OSPF Redistribution Mapping Dialog Box 67-39
OSPF Max Prefix Mapping Dialog Box 67-41
RIP Routing on Cisco IOS Routers 67-42
Defining RIP Setup Parameters 67-42
Defining RIP Interface Authentication Settings 67-43
Redistributing Routes into RIP 67-44
RIP Routing Policy Page 67-45
RIP Page—Setup Tab 67-45
RIP Page—Authentication Tab 67-46
RIP Authentication Dialog Box 67-47
RIP Page—Redistribution Tab 67-48
RIP Redistribution Mapping Dialog Box 67-49
Static Routing on Cisco IOS Routers 67-50
Defining Static Routes 67-50
Static Routing Policy Page 67-51
Static Routing Dialog Box 67-52

CHAPTER 68 **Managing Cisco Catalyst Switches and Cisco 7600 Series Routers** 68-1
Discovering Policies on Cisco Catalyst Switches and Cisco 7600 Series Routers 68-1
Viewing Catalyst Summary Information 68-2
Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups 68-3
Interfaces 68-5
Creating or Editing Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers 68-6
Deleting Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers 68-7
Interfaces/VLANs Page—Interfaces Tab 68-8
Create and Edit Interface Dialog Boxes—Access Port Mode 68-9
Create and Edit Interface Dialog Boxes—Routed Port Mode 68-12
Create and Edit Interface Dialog Boxes—Trunk Port Mode 68-14
Create and Edit Interface Dialog Boxes—Dynamic Mode 68-18
Create and Edit Interface Dialog Boxes—Subinterfaces 68-22
Create and Edit Interface Dialog Boxes—Unsupported Mode 68-24
VLANs 68-25
Creating or Editing VLANs 68-26
Deleting VLANs 68-27
Interfaces/VLANs Page—VLANs Tab 68-27

Create and Edit VLAN Dialog Boxes 68-28
Access Port Selector Dialog Box 68-30
Trunk Port Selector Dialog Box 68-31
VLAN Groups 68-31
Creating or Editing VLAN Groups 68-32
Deleting VLAN Groups 68-33
Interfaces/VLANs Page—VLAN Groups Tab 68-33
Create and Edit VLAN Group Dialog Boxes 68-34
Service Module Slot Selector Dialog Box 68-35
VLAN Selector Dialog Box 68-35
VLAN ACLs (VACLs) 68-36
Creating or Editing VACLs 68-37
Deleting VACLs 68-38
VLAN Access Lists Page 68-39
Create and Edit VLAN ACL Dialog Boxes 68-41
Create and Edit VLAN ACL Content Dialog Boxes 68-41
IDSM Settings 68-43
Creating or Editing EtherChannel VLAN Definitions 68-44
Deleting EtherChannel VLAN Definitions 68-45
Creating or Editing Data Port VLAN Definitions 68-46
Deleting Data Port VLAN Definitions 68-47
IDSM Settings Page 68-47
Create and Edit IDSM EtherChannel VLANs Dialog Boxes 68-49
Create and Edit IDSM Data Port VLANs Dialog Boxes 68-49

Monitoring, Reporting, and Diagnostics

CHAPTER 69 **Viewing Events** 69-1
Introduction to Event Viewer Capabilities 69-1
Historical View 69-2
Real-Time View 69-2
Views and Filters 69-3
Policy Navigation 69-3
Understanding Event Viewer Access Control 69-4
Scope and Limits of Event Viewer 69-4
Deeply Parsed Syslogs 69-6
Overview of Event Viewer 69-7
Event Viewer File Menu 69-9
Event Viewer View Menu 69-10

View List 69-12
Event Monitoring Window 69-14
Event Table Toolbar 69-16
Columns in Event Table 69-18
Time Slider 69-25
Event Details Pane 69-26
Preparing for Event Management 69-27
Ensuring Time Synchronization 69-27
Configuring ASA and FWSM Devices for Event Management 69-28
Configuring IPS Devices for Event Management 69-29
Managing the Event Manager Service 69-30
Starting, Stopping, and Configuring the Event Manager Service 69-30
Monitoring the Event Manager Service 69-31
Selecting Devices to Monitor 69-34
Monitoring Event Data Store Disk Space Usage 69-35
Archiving or Backing Up and Restoring the Event Data Store 69-36
Using Event Viewer 69-37
Using Event Views 69-37
Opening Views 69-38
Floating and Arranging Views 69-38
Customizing the Event Table Appearance 69-39
Switching Between Source/Destination IP Addresses and Host Object Names 69-39
Configuring Color Rules for a View 69-40
Creating Custom Views 69-41
Editing a Custom View Name or Description 69-41
Switching Between Real-Time and Historical Views 69-42
Saving Views 69-42
Deleting Custom Views 69-43
Filtering and Querying Events 69-43
Selecting the Time Range for Events 69-43
Using the Time Slider with Filtering 69-44
Refreshing the Event Table 69-44
Creating Column-Based Filters 69-45
Filtering Based on a Specific Event's Values 69-47
Filtering on a Text String 69-47
Clearing Filters 69-48
Performing Operations on Specific Events 69-49
Event Context (Right-Click) Menu 69-49
IPS Signature Quick Tune Dialog Box 69-52

Examining Details of a Single Event 69-53
Copying Event Records 69-53
Saving Events to a File 69-53
Looking Up a Security Manager Policy from Event Viewer 69-54
Looking Up Events for a Security Manager Policy 69-55
Viewing Events for an Access Rule 69-56
Viewing Events for an IPS Signature 69-57
Viewing Events for HPM Devices and Site-to-Site VPNs 69-58
Examples of Event Analysis 69-58
Help Desk: User Access To a Server Is Blocked By the Firewall 69-59
Monitoring and Mitigating Botnet Activity 69-61
Understanding the Syslog Messages That Indicate Actionable Events 69-61
Monitoring Botnet Using the Security Manager Event Viewer 69-62
Monitoring Botnet Using the Security Manager Report Manager 69-64
Monitoring Botnet Activity Using the Adaptive Security Device Manager (ASDM) 69-64
Mitigating Botnet Traffic 69-65
Removing False Positive IPS Events from the Event Table 69-66

CHAPTER 70 **Managing Reports** 70-1

Understanding Report Management 70-1
Understanding the Types of Reports Available in Security Manager 70-2
Preparing Devices for Report Manager Reporting 70-3
Understanding Report Manager Data Aggregation 70-4
Understanding Report Manager Access Control 70-5
Overview of Report Manager 70-6
Report Manager Menus 70-8
Understanding the Report List in Report Manager 70-9
Understanding the Report Settings Pane 70-10
Understanding the Generated Report Pane and Toolbar 70-12
Understanding the Predefined System Reports in Report Manager 70-13
Understanding Firewall Traffic Reports 70-14
Understanding Firewall Summary Botnet Reports 70-15
Understanding VPN Top Reports 70-16
Understanding General VPN Reports 70-16
Understanding IPS Top Reports 70-17
Understanding General IPS Reports 70-19
Working with Reports in Report Manager 70-19
Opening and Generating Reports 70-20
Creating Custom Reports 70-21

Editing Report Settings 70-22
Drilling Down into Report Data 70-26
Printing Reports 70-27
Exporting Reports 70-28
Configuring Default Settings for Reports 70-29
Arranging Report Windows 70-30
Saving Reports 70-31
Renaming Reports 70-31
Closing Report Windows 70-32
Deleting Reports 70-32
Managing Custom Reports 70-32
Scheduling Reports 70-33
Viewing Report Schedules 70-33
Configuring Report Schedules 70-34
Viewing Scheduled Report Results 70-35
Enabling and Disabling Report Schedules 70-36
Deleting Report Schedules 70-36
Troubleshooting Report Manager 70-36

CHAPTER 71 **Health and Performance Monitoring** 71-1
Health and Performance Monitor Overview 71-1
Trend Information 71-2
Monitoring Multiple Contexts 71-3
HPM Access Control 71-3
Preparing for Health and Performance Monitoring 71-4
Launching the Health and Performance Monitor 71-4
Managing Monitored Devices 71-5
HPM Window 71-6
Working with Table Columns 71-8
Showing and Hiding Table Columns 71-8
Column-based Filtering 71-17
Using The List Filter Fields 71-19
Monitoring Devices 71-21
Managing Device Views 71-21
Views: Opening and Closing 71-23
Views: Tiling Horizontally or Vertically 71-23
Views: Floating and Docking 71-24
Views: Custom 71-24
HPM Window: Monitoring Display 71-25

Monitoring Views: Devices or VPNs Summary 71-27
Monitoring Views: Device or VPN Status List 71-27
Monitoring Views: Device or VPN Details 71-28
Monitoring Views: VPN, RA and S2S 71-30
Exporting HPM Data 71-31
Alerts and Notifications 71-32
HPM Window: Alerts Display 71-32
Alerts: Configuring 71-34
Alerts Configuration: IPS 71-35
Alerts Configuration: Firewall 71-37
Alerts Configuration: VPN 71-39
Alerts: Viewing 71-41
Alerts: Acknowledging and Clearing 71-42
Alerts: History 71-43
SNMP Trap Forwarding Notification 71-44
SNMP Trap Entries Dialog Box 71-45
Add/Edit/Copy SNMP Trap Entries Dialog Box 71-46

CHAPTER 72 **Using External Monitoring, Troubleshooting, and Diagnostic Tools** 72-1
Dashboard Overview 72-1
CSM Mobile 72-11
Viewing Inventory Status 72-12
Inventory Status Window 72-13
Starting Device Managers 72-14
Troubleshooting Device Managers 72-16
Access Rule Look-up from Device Managers 72-17
Navigating to an Access Rule from ASDM 72-18
Navigating to an Access Rule from SDM 72-19
Launching Cisco Prime Security Manager or FireSIGHT Management Center 72-20
Detecting ASA CX and FirePOWER Modules 72-21
Sharing Device Inventory and Policy Objects with PRSM 72-22
Analyzing an ASA or PIX Configuration Using Packet Tracer 72-23
Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools 72-26
Analyzing Configuration Using Ping 72-26
Analyzing Configuration Using TraceRoute 72-28
Analyzing Configuration Using NS Lookup 72-29
Using the Packet Capture Wizard 72-30
IP Intelligence 72-33

Integrating CS-MARS and Security Manager 72-36
Checklist for Integrating CS-MARS with Security Manager 72-36
Configuring the Security Manager Server to Respond to CS-MARS Policy Queries 72-37
Registering CS-MARS Servers in Security Manager 72-38
Discovering or Changing the CS-MARS Controllers for a Device 72-39
Troubleshooting Tips for CS-MARS Querying 72-40
Looking Up CS-MARS Events for a Security Manager Policy 72-41
Viewing CS-MARS Events for an Access Rule 72-42
Viewing CS-MARS Events for an IPS Signature 72-44
Looking Up a Security Manager Policy from a CS-MARS Event 72-45
System Log Messages Supported for Policy Look-up 72-46
NetFlow Event Reporting in CS-MARS 72-47

Image Management

CHAPTER 73 **Using Image Manager 73-1**
Getting Started with Image Manager 73-1
Image Manager Supported Platforms and Versions 73-2
Device Configurations supported by Image Manager 73-4
Image Management for Multi-Context ASA 73-5
Image Manager Supported Image Types 73-5
Administrative Settings for Image Manager 73-6
Bootstrapping Devices for Image Manager 73-8
Working with Images 73-9
View All Images 73-10
Download Images to the Repository 73-11
Working with Bundles 73-13
Creating Bundles 73-13
View Images by Bundle 73-14
Renaming Bundles 73-15
Deleting Bundles 73-15
Deleting Images from Bundles 73-15
Working with Devices 73-16
Viewing Device Inventory 73-16
Manage Images on a Device 73-17
View Device Memory 73-18
Configuring the Image Install Location 73-19
About Image Updates on Devices Using Image Manager 73-20
Validating a Proposed Image Update on a Device 73-23

Using the Image Installation Wizard to Install Images on Devices 73-26
Install Bundled Images on Devices 73-30
Install Compatible Images on Devices 73-30
Install Images on Selected Devices 73-31
Working with Jobs 73-32
Viewing Image Installation Job Summary 73-33
Viewing Install Jobs 73-34
Aborting an Image Installation Job 73-35
Retry a Failed Image Install Job 73-35
Roll Back a Deployed Job 73-35
Image Installation Job Approval Workflow 73-36
Troubleshooting Image Management 73-37

INDEX

Preface

Conventions

This document uses the following conventions:

Item	Convention
Commands, keywords, special terminology, and options that should be selected during procedures	**boldface** font
Variables for which you supply values and new or important terminology	*italic* font
Displayed session and system information, paths and file names	`screen` font
Information you enter	**boldface screen** font
Variables you enter	*italic screen* font
Menu items and button names	**boldface** font
Indicates menu items to select, in the order you select them.	**Option > Network Preferences**

Tip Identifies information to help you get the most benefit from your product.

Note Means *reader take note*. Notes identify important information that you should reflect upon before continuing, contain helpful suggestions, or provide references to materials not contained in the document.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage, loss of data, or a potential breach in your network security.

Warning **Identifies information that you must heed to prevent damaging yourself, the state of software, or equipment. Warnings identify definite security breaches that will result if the information presented is not followed carefully.**

Obtain Documentation and Submit a Service Request

For information on obtaining documentation, using the Cisco Bug Search Tool (BST), submitting a service request, and gathering additional information, see *What's New in Cisco Product Documentation*.

To receive new and revised Cisco technical content directly to your desktop, you can subscribe to the What's New in Cisco Product Documentation RSS feed. The RSS feeds are a free service.

PART 1

The Basics of Using Security Manager

Getting Started with Security Manager

The following topics describe Cisco Security Manager, how to get started with the application, and how to complete its configuration.

- Product Overview, page 1-1
- Logging In to and Exiting Security Manager, page 1-11
- Using Configuration Manager - Overview, page 1-14
- Using the JumpStart to Learn About Security Manager, page 1-24
- Completing the Initial Security Manager Configuration, page 1-25
- Understanding Basic Security Manager Interface Features, page 1-29
- Accessing Online Help, page 1-52

Product Overview

Note From version 4.17, though Cisco Security Manager continues to manage the following devices, it will not provide support for any enhancements:

- Cisco Catalyst 6500 and 7600 Series Firewall Services Modules (EOL8184)
- Cisco Catalyst 6500 Series Intrusion Detection System Services Module 2 (EOL8843)
- Cisco Intrusion Prevention System: IPS 4200, 4300, and 4500 Series Sensors (EOL9916)
- Cisco SR 500 Series Secure Routers (EOL7687, EOL7657)
- PIX Firewalls (EOL)

Caution From version 4.18, Cisco Security Manager does not support SFR from ASA 9.10(1) onwards for ASA 5512, ASA 5506, ASA 5506H and ASA 5506W models. Therefore, if you upgrade to 9.10(1) through Image Manager, the exiting SFR configuration will be lost.

Cisco Security Manager (Security Manager) enables you to manage security policies on Cisco security devices. Security Manager supports integrated provisioning of firewall, and VPN (site-to-site, remote access, and SSL) services across ASA security appliances.

For a complete list of devices and OS versions supported by Security Manager, please refer to *Supported Devices and Software Versions for Cisco Security Manager* on Cisco.com.

Security Manager also supports provisioning of many platform-specific settings, for example, interfaces, routing, identity, QoS, logging, and so on.

Security Manager efficiently manages a wide range of networks, from small networks consisting of a few devices to large networks with thousands of devices. Scalability is achieved through a rich feature set of shareable objects and policies and device grouping capabilities.

Security Manager supports multiple configuration views optimized around different task flows and use cases.

The following topics provide an overview of Security Manager:

- Primary Benefits of Cisco Security Manager, page 1-2
- Security Manager Policy Feature Sets, page 1-4
- Security Manager Applications Overview, page 1-6
- Device Monitoring Overview, page 1-7
- IPv6 Support in Security Manager, page 1-8

Primary Benefits of Cisco Security Manager

These are the primary benefits of working with Security Manager:

- **Scalable network management**—Centrally administer security policies and device settings for either small networks or large scale networks consisting of thousands of devices. Define policies and settings once and then optionally assign them to individual devices, groups of devices or all the devices in the enterprise.
- **Provisioning of multiple security technologies across different platforms**—Manage VPN, firewall, and IPS technologies on routers, security appliances, Catalyst devices and service modules, and IPS devices.
- **Provisioning of platform-specific settings and policies**—Manage platform-specific settings on specific device types. For example: routing, 802.1x, EzSDD, and Network Admission Control on routers, and device access security, DHCP, AAA, and multicast on firewall devices.
- **VPN wizards**—Quickly and easily configure point-to-point, hub-and-spoke, full-mesh, and Extranet site-to-site VPNs across different VPN device types. Quickly and easily configure remote access IPsec and SSL VPNs on ASA, IOS, and PIX devices.
- **Multiple management views**—Device, policy, and map views enable you to manage your security in the environment that best suits your needs.
- **Reusable policy objects**—Create reusable objects to represent network addresses, device settings, VPN parameters, and so on, then use them instead of manually entering values.
- **Device grouping capabilities**—Create device groups to represent your organizational structure. Manage all devices in the groups concurrently.
- **Policy inheritance**—Centrally specify which policies are mandatory and enforced lower in the organization.
- **Role-based administration**—Enable appropriate access controls for different operators.
- **Workflow**—Optionally allow division of responsibility and workload between network operators and security operators and provide a change management approval and tracking mechanism.

- **Ticket Management**—Associate a ticket ID with policy changes, easily add and update comments pertaining to those changes, and quickly navigate to an external change management system from Security Manager.
- **Single, consistent user interface for managing common firewall features**—Single rule table for all platforms (router, PIX, ASA, and FWSM).
- **Image management**—Complete image management for ASA devices. Facilitates at every stage of image upgrade of devices by: downloading and maintaining image repository, evaluating images, analyzing impact of upgrades, preparing and planning reliable and stable device upgrades, and ensuring sufficient fallback and recovery mechanisms.
- **Intelligent analysis of firewall policies**—The conflict detection feature analyzes and reports rules that overlap or conflict with other rules. The ACL hit count feature checks in real-time whether specific rules are being hit or triggered by packets.
- **Sophisticated rule table editing**—In-line editing, ability to cut, copy, and paste rules and to change their order in the rule table.
- **Discover firewall policies from device**—Policies that exist on the device can be imported into Security Manager for future management.
- **Flexible deployment options**—Support for deployment of configurations directly to a device or to a configuration file. You can also use Auto-Update Server (AUS), Configuration Engine, or Token Management Server (TMS) for deployment.
- **Rollback**—Ability to roll back to a previous configuration if necessary.
- **FlexConfig (template manager)**—Intelligent CLI configlet editor to manage features available on a device but not natively supported by Security Manager.
- **Integrated device monitoring and reporting**—Features for monitoring events on IPS, ASA, and FWSM devices and correlating them to the related configuration policies, and for creating security and usage reports. These features include the following stand-alone Security Manager applications:
  - **Event Viewer**—Event Viewer monitors your network for system log (syslog) events from ASA and FWSM devices, as well as security contexts and SDEE events from IPS devices and virtual sensors. Event Viewer collects these events and provides an interface by which you can view them, group them, and examine their details in near real time.
  - **Report Manager**—Report Manager lets you collect, display and export a wide variety of network usage and security information for ASA and IPS devices, and for ASA-hosted remote-access IPsec and SSL VPNs. These reports aggregate security data such as top sources, destinations, attackers, victims, as well as security information such as top bandwidth, duration, and throughput users. Data is available for hourly, daily, and monthly periods. (Report Manager aggregates information collected from devices monitored by the Event Manager service. Thus, to view reports about a device, you must be monitoring that device in Event Viewer.)

    Note Report Manager does not report FWSM events even though Event Viewer works with FWSM.

  - **Health and Performance Monitor**—Health and Performance Monitor (HPM) periodically polls monitored ASA devices, IPS devices, and ASA-hosted VPN services for key health and performance data, including critical and non-critical issues, such as memory usage, interface status, dropped packets, tunnel status, and so on. This information is used for alert generation and email notification, and to display trends based on aggregated data, which is available for hourly, daily, and weekly periods.

Note Health and Performance Monitor does not monitor FWSM devices.

- **Dashboard**—The Dashboard is a configurable launch point for Security Manager that makes IPS and FW tasks more convenient for you. In addition to the original dashboard, you can create new, additional dashboards, and you can customize all dashboards. By using the dashboard, you can accomplish in one place many tasks that are found in several other areas of Security Manager, such as the IPS Health Monitor page, Report Manager, Health and Performance Monitor, and IP Intelligence Settings. For detailed information on the dashboard, see Dashboard Overview, page 72-1.

Additional features let you monitor devices from Security Manager using other closely related applications, including Cisco Security Monitoring, Analysis and Response System (CS-MARS), Cisco Performance Monitor, and device managers such as ASDM (read-only versions of which are included with Security Manager).

Security Manager Policy Feature Sets

Security Manager provides the following primary feature sets for configuration policies:

- **Firewall Services**

  Configuration and management of firewall policies across multiple platforms, including IOS routers, ASA/PIX devices, and Catalyst Firewall Service Modules (FWSMs). Features include:

  - Access control rules—Permit or deny traffic on interfaces through the use of access control lists for both IPv4 and IPv6 traffic.
  - Botnet Traffic Filter rules—(ASA only.) Filter traffic based on known malware sites and optionally drop traffic based on threat level.
  - Inspection rules—Filter TCP and UDP packets based on application-layer protocol session information.
  - AAA/Authentication Proxy rules—Filter traffic based on authentication and authorization for users who log into the network or access the Internet through HTTP, HTTPS, FTP, or Telnet sessions.
  - Web filtering rules—Use URL filtering software, such as Websense, to deny access to specific web sites.
  - ScanSafe Web Security—(Routers only.) Redirect HTTP/HTTPS traffic to the ScanSafe web security center for content scanning and malware protection services.
  - Transparent firewall rules—Filter layer-2 traffic on transparent or bridged interfaces.
  - Zone-based firewall rules—Configure access, inspection, and web filtering rules based on zones rather than on individual interfaces.

  For more information, see Chapter 12, “Introduction to Firewall Services”.

- **Site-to-Site VPN**

  Setup and configuration of IPsec site-to-site VPNs. Multiple device types can participate in a single VPN, including IOS routers, PIX/ASA devices, and Catalyst VPN Service Modules. Supported VPN topologies are:

  - Point to point
  - Hub and spoke

  - Full mesh
  - Extranet (a point-to-point connection to an unmanaged device)

  Supported IPsec technologies are:

  - Regular IPsec
  - GRE
  - GRE Dynamic IP
  - DMVPN
  - Easy VPN
  - GET VPN

  For more information, see Chapter 25, "Managing Site-to-Site VPNs: The Basics".

- **Remote Access VPN**

  Setup and configuration of IPsec and SSL VPNs between servers and mobile remote workstations running Cisco VPN client or AnyConnect client software. For more information, see Chapter 30, "Managing Remote Access VPNs: The Basics".

- **Intrusion Prevention System (IPS) Management**

  Management and configuration of Cisco IPS sensors (appliances and service modules) and IOS IPS devices (Cisco IOS routers with IPS-enabled images and Cisco Integrated Services Routers).

  For more information, see Overview of IPS Configuration, page 36-5 and Overview of Cisco IOS IPS Configuration, page 45-4.

- **Features Specific to Firewall Devices (PIX/ASA/FWSM)**

  Configuration of advanced platform-specific features and settings on PIX/ASA devices and Catalyst FWSMs. These features provide added value when managing security profiles and include:

  - Interface configuration
  - Identity-aware firewall settings
  - Device administration settings
  - Security
  - Routing
  - Multicast
  - Logging
  - NAT
  - Bridging
  - Failover
  - Security contexts

  For more information, see Chapter 46, "Managing Firewall Devices".

- **Features Specific to IOS Routers**

  Configuration of advanced platform-specific features and settings on IOS routers. These features provide added value when managing security profiles and include:

  - Interface configuration
  - Routing

  - NAT
  - 802.1x
  - NAC
  - QoS
  - Dialer interfaces
  - Secure device provisioning

  For more information, see Chapter 61, "Managing Routers".

- **Features Specific to Catalyst 6500/7600 Devices and Catalyst Switches**

  Configuration of VLAN, network connectivity, and service module features and settings on Catalyst 6500/7600 devices and on other Catalyst switches.

  For more information, Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers".

- **FlexConfigs**

  Flexconfig policies and policy objects enable you to provision features that are available on the device but not natively supported by Security Manager. They enable you to manually specify a set of CLI commands and to deploy them to devices using Security Manager's provisioning mechanisms. These commands can be either prepended or appended to the commands generated by Security Manager to provision security policies.

  For more information, see Chapter 7, "Managing FlexConfigs".

Security Manager Applications Overview

The Security Manager client has six main applications and one application designed for mobile devices:

- **Configuration Manager**—This is the primary application. You use Configuration Manager to manage the device inventory, create and edit local and shared policies, manage VPN configurations, and deploy policies to devices. Configuration Manager is the largest of the applications and most of the documentation addresses this application. If a procedure does not specifically mention an application, the procedure is using Configuration Manager. For an introduction to Configuration Manager, see Using Configuration Manager - Overview, page 1-14.
- **Event Viewer**—This is an event monitoring application, where you can view and analyze events generated from IPS, ASA, and FWSM devices that you have configured to send events to Security Manager. For information about using Event Viewer, see Chapter 69, "Viewing Events".
- **Report Manager**—This is a reporting application, where you can view and create reports of aggregated information on device and VPN statistics. Much of the information is derived from events available through Event Viewer, but some of the VPN statistics are obtained by communicating directly with the device. For information about using Report Manager, see Chapter 70, "Managing Reports".
- **Health & Performance Monitor**—The HPM application lets you monitor key health and performance data for ASA (including ASA-SM) devices, IPS devices, and VPN services by providing network-level visibility into device status and traffic information. This ability to monitor key network and device metrics lets you quickly detect and resolve device malfunctions and bottlenecks in the network. See Chapter 71, "Health and Performance Monitoring" for more information about this application.

- **Image Manager**—The Image Manager application provides complete image management of ASA devices. It facilitates downloading, evaluating, analyzing, preparing, and planning image updates. It assesses image availability, compatibility, and impact on devices and provides scheduling, grouping, and change management of device updates. In addition, Image Manager includes capabilities for maintaining an image repository as well as for ensuring stable fallback and recovery mechanisms for image updates on ASA devices. For information about using Image Manager, see Chapter 73, "Using Image Manager".
- **Dashboard**—The Dashboard is a configurable launch point for Security Manager that makes IPS and FW tasks more convenient for you. In addition to the original dashboard, you can create new, additional dashboards, and you can customize all dashboards. By using the dashboard, you can accomplish in one place many tasks that are found in several other areas of Security Manager, such as the IPS Health Monitor page, Report Manager, Health and Performance Monitor, and IP Intelligence Settings. For detailed information on the dashboard, see Dashboard Overview, page 72-1.

You can open any of these applications directly from the Windows Start menu or a desktop icon, or you can open them from within any of these applications through the application's Launch menu. For information on opening applications, see Logging In to and Exiting the Security Manager Client, page 1-12.

The Security Manager client has an additional application, CSM Mobile, which is designed specifically for mobile devices:

- **CSM Mobile**—CSM Mobile allows you to access device health summary information from mobile devices. The information available to you in this way is the same as that available in the Device Health Summary widget in the Dashboard: current high or medium severity active alerts generated by HPM. Alerts can be grouped by Alert-Description, Predefined-Category, Device, or Alert Technology. For more information on CSM Mobile, see CSM Mobile, page 72-11. For more details on device health summary information, see Dashboard Overview, page 72-1. For information on enabling or disabling CSM Mobile, see CSM Mobile Page, page 11-9

Device Monitoring Overview

Security Manager includes several facilities for monitoring devices:

- **Event Viewer**—This integrated tool allows you to view events on ASA, FWSM, and IPS devices and correlate them to the related configuration policies. This helps you identify problems, troubleshoot configurations, and then fix the configurations and redeploy them. For more information, see Chapter 69, "Viewing Events".
- **Report Manager**—This is a reporting application, where you can view and create reports of aggregated information on device and VPN statistics. Much of the information is derived from events available through Event Viewer, but some of the VPN statistics are obtained by communicating directly with the device. For information about using Report Manager, see Chapter 70, "Managing Reports".

  For information on all of the types of reports available in Security Manager, see Understanding the Types of Reports Available in Security Manager, page 70-2.
- **Health & Performance Monitor**—The HPM application lets you monitor key health and performance data for ASA (including ASA-SM) devices, IPS devices, and VPN services by providing network-level visibility into device status and traffic information. See Chapter 71, "Health and Performance Monitoring" for more information about this application.

- **Dashboard**—The Dashboard is a configurable launch point for Security Manager that makes IPS and FW tasks more convenient for you. In addition to the original dashboard, you can create new, additional dashboards, and you can customize all dashboards. By using the dashboard, you can accomplish in one place many tasks that are found in several other areas of Security Manager, such as the IPS Health Monitor page, Report Manager, Health and Performance Monitor, and IP Intelligence Settings. For detailed information on the dashboard, see Dashboard Overview, page 72-1.
- **Packet Tracer**—You can use this tool to test whether certain types of packets will be allowed to go through an ASA device. For more information, see Analyzing an ASA or PIX Configuration Using Packet Tracer, page 72-23.
- **Ping, Trace route, and NS Lookup**—You can use ping and traceroute on a managed device to check whether there is a route between the device and a specific destination. You can use NS lookup to resolve addresses to DNS names. For more information, see Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools, page 72-26.
- **Cisco Prime Security Manager (PRSM) Integration**—You can "cross launch" PRSM from the Configuration Manager application. The PRSM application is used to configure and manage ASA CX devices. For more information, see Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20.
- **Device Manager Integration**—Security Manager includes read-only copies of the various device managers, such as Adaptive Security Device Manager (ASDM). You can use these tools to view device status, but not to change the device configuration. For more information, see Starting Device Managers, page 72-14.
- **Cisco Security Monitoring, Analysis and Response System (CS-MARS) Integration**—If you use the CS-MARS application, you can integrate it with Security Manager and view events in CS-MARS from Security Manager, and conversely, Security Manager policies related to events from CS-MARS. For more information, see Integrating CS-MARS and Security Manager, page 72-36.

IPv6 Support in Security Manager

Security Manager provides increasing support for IPv6 configuration, monitoring, and reporting.

Beginning with version 4.12, Security Manager supports communication from Security Manager server to the managed devices over either IPv6 address or IPv4 address. This feature is available only for firewall devices, that is, those devices where the OS type is either ASA or FWSM. To enable communication over IPv6 addresses, you must first enable IPv6 address on the Security Manager server. See Configuring IPv6 on Security Manager Server, page 1-9 for more information.

Note The communication between Security Manager server and Security Manager client is over IPv4 address only. IPv6 address is not supported for server to client communication. Also, if ACS server is used for authentication, the ACS must have IPv4 address. IPv6 communication to ACS server is not supported. Auto Update Server (AUS) does not support IPv6 addresses.

For versions prior to 4.12, to manage a device that supports IPv6 addressing with Security Manager, you must configure the device's management address as an IPv4 address. All communications between the device and Security Manager, such as policy discovery and deployment, use IPv4 transport. If the IPv6 policies are not appearing for a supported device, rediscover the device policies; if necessary, delete the device from the inventory and add it again.

Configuring IPv6 on Security Manager Server

Follow these steps to configure IPv6 on Security Manager server for communicating with a device over IPv6 address.

Step 1 On the Security Manager server, go to **Start > Control panel > Network and Internet > Network Connections**.

Step 2 Click the available Network Connection to open the **Ethernet Status** window. Click **Properties**. The Ethernet Properties window appears.

Step 3 On the Networking tab, check the **Internet Protocol Version 6 (TCP/IPv6)** check box, and then click **Properties**. The Internet Protocol Version 6 (TCP/IPv6) Properties window appears.

Step 4 Configure the IPv6 static address and DNS servers, and click **OK**.

Note You must configure Security Manager server hostname to resolve to IPv4 addresses only. The server hostname should not resolve to IPv6 address.

Configuring IPv6 Policies

In general, you can configure IPv6 policies on the following types of device. In addition, you can monitor IPv6 alerts generated by IPS, ASA, and FWSM devices. For other types of devices, use FlexConfig policies to configure IPv6 settings. For more specific information on IPv6 device support, see the *Supported Devices and Software Versions for Cisco Security Manager* document on Cisco.com.

- **ASA**—Release 7.0+ when running in router mode; release 8.2+ when running in transparent mode. Both single and multiple security context devices are supported.
- **FWSM**—Release 3.1+ when running in router mode. Not supported in transparent mode. Both single and multiple security context devices are supported.
- **IPS**—Release 6.1+.

Following is a summary of the Security Manager features that support IPv6 addressing:

- **Policy Objects**—The following policy objects support IPv6 addresses:
  - Networks/Hosts. See Understanding Networks/Hosts Objects, page 6-80.
  - Services. This object includes predefined services for ICMP6 and DHCPv6, which you can use only with IPv6 policies. The other services apply to both IPv4 and IPv6. For more information on service objects, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.
- **Firewall Services Policies**—The following Firewall Services policies and tools support IPv6 configurations:
  - AAA Rules. See Chapter 15, “Managing Firewall AAA Rules”.
  - Access Rules. See Configuring Access Rules, page 16-7.
  - Inspection Rules. See Chapter 17, “Managing Firewall Inspection Rules”.
  - Settings > Access Control. See Configuring Settings for Access Control, page 16-23.
  - Tools:

    Hit Count. See Viewing Hit Count Details, page 16-36.

    Find and Replace. See Finding and Replacing Items in Rules Tables, page 12-16.

- **ASA and FWSM Policies**—The following ASA and FWSM policies support IPv6 configurations:
  - (ASA 7.0+ routed mode; ASA 8.2+ transparent mode; FWSM 3.1+ routed mode.) Interfaces: IPv6 tab of the Add Interface and Edit Interface dialog boxes. See Configuring IPv6 Interfaces (ASA/FWSM), page 46-47.
  - (ASA only.) Platform > Bridging > IPv6 Neighbor Cache. See Managing the IPv6 Neighbor Cache, page 47-7.
  - (ASA 5505 8.2/8.3 only.) Platform > Bridging > Management IPv6. See Management IPv6 Page (ASA 5505), page 47-11.
  - (ASA 8.4.2+ only.) Platform > Device Admin > Server Access > DNS. See DNS Page, page 52-14.
- **FlexConfig Policies**—There are two Firewall system variables that you can use to identify IPv6 ACLs on a device. For more information, see FlexConfig System Variables, page 7-7.

  There is also a predefined FlexConfig policy object that uses these variables, ASA_add_IPv6_ACEs.
- **Event Viewer**—Events that include IPv6 addresses are supported, and the addresses are displayed in the same columns as IPv4 addresses: Source, Destination, and IPLog Address (for IPS alerts). However, you must configure the device to use IPv4 for sending events to the Security Manager server. All event communications use IPv4 transport. For more information on Event Viewer, see Chapter 69, “Viewing Events”.
- **Dashboard**—On the Dashboard, all the widgets that use IP addressing support IPv6 addresses. However, as is true elsewhere in Security Manager, you must configure the device to use IPv4 for sending events to the Security Manager server. All event communications use IPv4 transport. For more information on the Dashboard, see Dashboard Overview, page 72-1
- **Report Manager**—Reports include statistics for IPv6 events collected by Event Management. For more information on Report Manager, see Chapter 70, “Managing Reports”.

Policy Object Changes in Security Manager 4.4

Certain changes were made to a few policies and policy objects in Security Manager 4.4, in order to unify previously separate IPv4 and IPv6 elements. The most important of these changes are to the Networks/Hosts object (which itself represents a unification of the Networks/Hosts and the Networks/Hosts-IPv6 objects):

- The new Networks/Hosts object “All-IPv4-Addresses” replaces the IPv4 “any” network policy object. If you upgrade to Security Manager 4.4 from a previous version, all references to the IPv4 “any” network policy object will be changed to “All-IPv4-Addresses.”
- The new Networks/Hosts object “All-IPv6-Addresses” replaces the IPv6 “any” network policy object. If you upgrade to Security Manager 4.4 from a previous version, all references to the IPv6 “any” network policy object will be changed to “All-IPv6-Addresses.”
- The new Networks/Hosts object “All-Addresses” does not have a corresponding policy object in earlier versions of Security Manager. It is a new global “any” policy object, and it encompasses all IPv4 and IPv6 address ranges.

Other related changes include unification of IPv4 and IPv6 versions of device-specific policies such as Access Rules, Inspection Rules, and so on.

Further, when editing policies and objects, IPv4, IPv6, or mixed-mode (both IPv4 and IPv6) entries are automatically filtered in elements, such as dialog boxes, in which one or more of those entries is not appropriate to that element.

Related Topics

- Policy Object Manager, page 6-4
- Understanding Networks/Hosts Objects, page 6-80

Logging In to and Exiting Security Manager

Security Manager has two main interfaces:

- Cisco Security Management Suite home page—Use this interface to install the Security Manager client and to manage the server. You can also access other CiscoWorks applications you installed, such as Resource Manager Essentials (RME).
- Security Manager clients—Use these interfaces to perform most Security Manager tasks. You can log directly into any of six client applications: Configuration Manager, Event Viewer, Report Manager, Health & Performance Monitor, Image Manager, and Dashboard.

These topics describe how to log in to and exit these interfaces:

- Understanding User Permissions, page 1-11
- Logging In to the Cisco Security Management Suite Server, page 1-12
- Logging In to and Exiting the Security Manager Client, page 1-12

Understanding User Permissions

Cisco Security Manager authenticates your username and password before you can log in. After you are authenticated, Security Manager establishes your role within the application. This role defines your permissions (also called privileges), which are the set of tasks or operations that you are authorized to perform. If you are not authorized for certain tasks or devices, the related menu items, items in tables of contents, and buttons are hidden or disabled. In addition, a message tells you that you do not have permission to view the selected information or perform the selected operation.

Authentication and authorization for Security Manager is managed either by the CiscoWorks server or the Cisco Secure Access Control Server (ACS). By default, CiscoWorks manages authentication and authorization, but you can configure Security Manager to use your Cisco Secure ACS setup.

When using ACS, if all of the ACS servers become unavailable, you cannot perform tasks in Security Manager. If you are logged in, you might be abruptly logged out of the system (without an opportunity to save changes) if you try to perform a task that requires ACS authorization. If this happens, you get a message stating this is the reason you are getting logged off.

For more information about user permissions and AAA configuration, see the *Installation Guide for Cisco Security Manager*.

For more information about authorization control in the Event Viewer and Report Manager applications, see the following topics:

- Understanding Event Viewer Access Control, page 69-4
- Understanding Report Manager Access Control, page 70-5

Logging In to the Cisco Security Management Suite Server

Use the Cisco Security Management Suite home page, and CiscoWorks Common Services, to install the Security Manager client and to manage the server. You can also access other CiscoWorks applications you installed, such as RME.

Note The **Software Center > Software Update** feature in Common Services is not supported by Cisco Security Manager.

Step 1 In your web browser, open one of these URLs, where *SecManServer* is the name of the computer where Security Manager is installed. Click **Yes** on any Security Alert windows.

- If you are not using SSL, open http://*SecManServer*:1741
- If you are using SSL, open https://*SecManServer*:443

The Cisco Security Management Suite login screen is displayed. Verify on the page that JavaScript and cookies are enabled and that you are running a supported version of the web browser. For information on configuring the browser to run Security Manager, see *Installation Guide for Cisco Security Manager*.

Step 2 Log in to the Cisco Security Management Suite server with your username and password. When you initially install the server, you can log in using the username **admin** and the password defined during product installation.

Step 3 On the Cisco Security Management Suite home page, you can access at least the following features. Other features might be available depending on how you installed the product.

- Cisco Security Manager Client Installer—Click this item to install the Security Manager client. The client is the main interface for using the product.
- Server Administration—Click this item to open the CiscoWorks Common Services Server page. CiscoWorks Common Services is the foundation software that manages the server. Use it to configure and manage back-end server features such as server maintenance and troubleshooting, local user definition, and so on.
- CiscoWorks link (in the upper right of the page)—Click this link to open the CiscoWorks Common Services home page.

Step 4 To exit the application, click **Logout** in the upper right corner of the screen. If you have both the home page and the Security Manager client open at the same time, exiting the browser connection does not exit the Security Manager client.

Note To meet PCI compliance, TLS 1.0 is disabled from CSM server. Hence, CSM server will not allow any TLS 1.0 clients to connect. This change is not applicable for CSM server to device communication. Existing CSM server to device communication will be supported as is.

Logging In to and Exiting the Security Manager Client

Use the Security Manager client to perform most Security Manager tasks.

Tip You must log into the workstation using a Windows user account that has Administrator privileges to fully use the Security Manager client applications. If you try to operate the applications with lesser privileges, you might find that some features do not work correctly.

Before You Begin

Install the client on your computer. To install the client, log into the Security Manager server as described in Logging In to the Cisco Security Management Suite Server, page 1-12, and then click **Cisco Security Manager Client Installer** and follow the instructions in the installation wizard.

Step 1 Select one of the following applications from the **Start > All Programs > Cisco Security Manager Client** menu:

- Configuration Manager
- Event Viewer
- Report Manager
- Health & Performance Monitor
- Image Manager
- Dashboard

Tip If the client was installed on the workstation, but it does not appear in your Start menu, it probably was installed by another user. To make Security Manager Client visible in the Start menu for every user of the client station, copy the Cisco Security Manager Client folder from Documents and Settings\<user>\Start Menu\Programs\Cisco Security Manager to Documents and Settings\All Users\Start Menu\Programs\Cisco Security Manager.

Step 2 In the application's login window, select the server to which you want to log in, and enter your Security Manager username and password. Click **Login**.

The client logs in to the server and opens the application you selected based on the following conditions. Note that these conditions are per application, for example, if you have Configuration Manager open on one workstation, opening Event Viewer from a different workstation has no implications for your Configuration Manager session unless or until you start Configuration Manager from Event Viewer.

- In both Workflow and non-Workflow mode, you cannot log into the same server from a single workstation and have more than one active session using the same user account. You are reminded that you are already logged in and asked to reuse the existing open application.
- In both workflow modes, you can log into different servers using the same (or different) user name from the same workstation.
- In non-Workflow mode, for a given server, if the user name is logged in on a different workstation, the client on the other workstation is automatically logged out, and any unsaved changes are lost. Thus, do not share user accounts, and if you must log in from different workstations to the same server, be sure to save your changes before leaving an active client.
- In Workflow mode, you can log in using the same user account multiple times but only from different workstations. However, you cannot open the same activity in Configuration Manager at the same time in more than one client; you must open different activities. Activities do not apply when using Event Viewer or Report Manager.

Tip The client automatically closes if it is idle for 120 minutes. To change the idle timeout, in Configuration Manager, select **Tools > Security Manager Administration**, select **Customize Desktop** from the table of contents, and enter the desired timeout period. You can also disable the feature so that the client does not close automatically. All applications use the same timeout setting, and working in one application resets the timer for all other applications.

Step 3 To exit the application, select **File > Exit**.

Using Configuration Manager - Overview

These topics provide an overview of the different views in which you can work in Configuration Manager, the basic task flow for defining and deploying policies to devices, and some basic concepts:

- Configuration Manager Overview, page 1-14
- Task Flow for Configuring Security Policies, page 1-19
- Policy and Policy Object Overview, page 1-20
- Workflow and Activities Overview, page 1-20

Configuration Manager Overview

The Configuration Manager application provides three views in which you can manage devices and policies: Device view, Policy view, and Map view. You can switch between these views according to your needs using toolbar buttons or the View menu.

- Device view—Provides a device-centric view, where you configure policies on specific devices. For more information, see Device View Overview, page 1-15.
- Policy view—Provides a policy-centric view, where you can create device-independent shared policies that you can assign to one or more devices. For more information, see Policy View Overview, page 1-16.
- Map view—Provides a visual representation of your network, which is primarily useful for visualizing and configuring site-to-site VPNs. For more information, see Map View Overview, page 1-18.

Each view presents a different way to access Configuration Manager functionality. What you can do, and how you do it, are determined by the view you select. In the Device and Policy views you see two selectors on the left and a work area on the right. In each of these, your selection in the upper selector determines what you can select in the lower selector. Your selection in the lower selector determines what you view in the work area. This design enables you to quickly and easily drill down to the network details that you want to view or edit.

Besides the main views, there are several additional tools used for configuring other items such as site-to-site VPNs and policy objects, or for monitoring devices. These tools are typically available from the Manage menu, although some are available on the Policy, Activities, Tools, or Launch menus. Some tools have related buttons in the toolbar. These tools open in a separate window so that you do not loose your place in the main view that you are currently using.

The following topics provide reference information about the basic features of the user interface:

- Menu Bar Reference for Configuration Manager, page 1-29
- Toolbar Reference (Configuration Manager), page 1-39
- Using Selectors, page 1-45
- Using Wizards, page 1-47
- Using Rules Tables, page 12-8
- Using Text Fields, page 1-49
- Accessing Online Help, page 1-52

Device View Overview

Device view in Configuration Manager enables you to add devices to the Security Manager inventory and to centrally manage device policies, properties, interfaces, and so on. The following figure identifies the functional areas of the Device view.

This is a device-centric view in which you can see all devices that you are managing and you can select specific devices to view their properties and define their settings and policies.

Note Security Manager also provides the ability to see the status of the devices in the Security Manager inventory. To access the Device Status View, select **View > Device Status View** or select one of the folder nodes in the Device selector. For more information, see Working with Device Status View, page 3-64.

In Device View, you can define security policies locally on specific devices. You can then share these policies to make them globally available to be assigned to other devices.

For more information, see Understanding the Device View, page 3-1.

Figure 1-1 Device View Overview

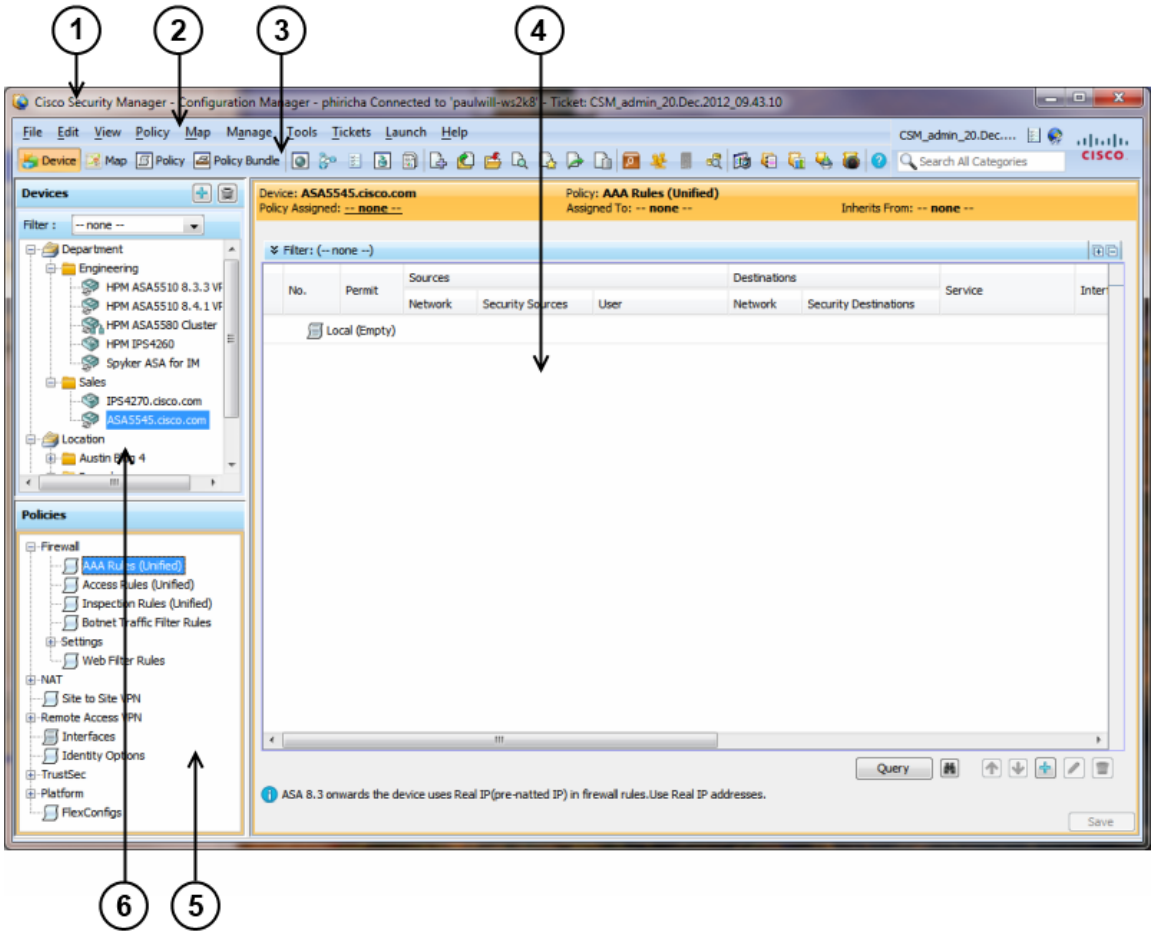

1	Title bar	**2**	Menu bar (see Menu Bar Reference for Configuration Manager, page 1-29)
3	Toolbar (see Toolbar Reference (Configuration Manager), page 1-39)	**4**	Work area
5	Policy selector	**6**	Device selector (see Using Selectors, page 1-45)

The title bar displays the following information about Security Manager:

- Your login name.
- The name of the Security Manager server to which you are connected.
- If Workflow mode is enabled, the name of the open activity.

Policy View Overview

Policy view in Configuration Manager enables you to create and manage reusable policies that can be shared among multiple devices. The following figure identifies the functional areas of the Policy view.

This is a policy-centric view in which you can see all the shareable policy types supported by Security Manager. You can select a specific policy type and create, view, or modify shared policies of that type. You can also see the devices to which each shared policy is assigned and change the assignments as required.

For more information, see Managing Shared Policies in Policy View, page 5-50.

Figure 1-2 Policy View Overview

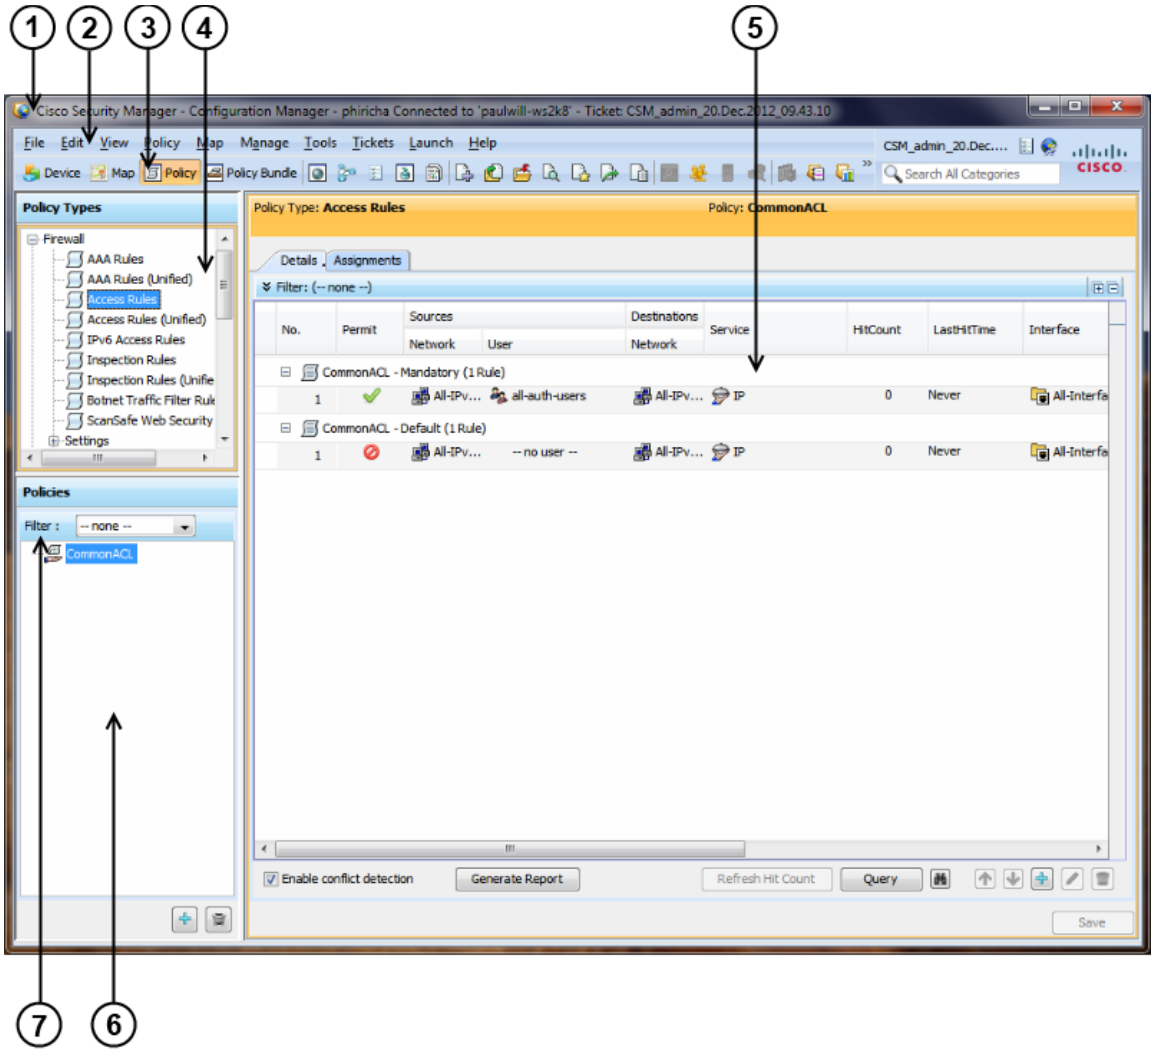

1	Title bar	**2**	Menu bar (see Menu Bar Reference for Configuration Manager, page 1-29)
3	Toolbar (see Toolbar Reference (Configuration Manager), page 1-39)	**4**	Policy type selector (see Using Selectors, page 1-45)
5	Work area	**6**	Shared policy selector
7	Policy filter		

Map View Overview

Map view in Configuration Manager enables you to create customized, visual topology maps of your network, within which you can view connections between your devices and easily configure VPNs and access control settings. The following figure identifies the functional areas of the Map view.

For more information, see Chapter 35, "Using Map View".

Figure 1-3 Map View Overview

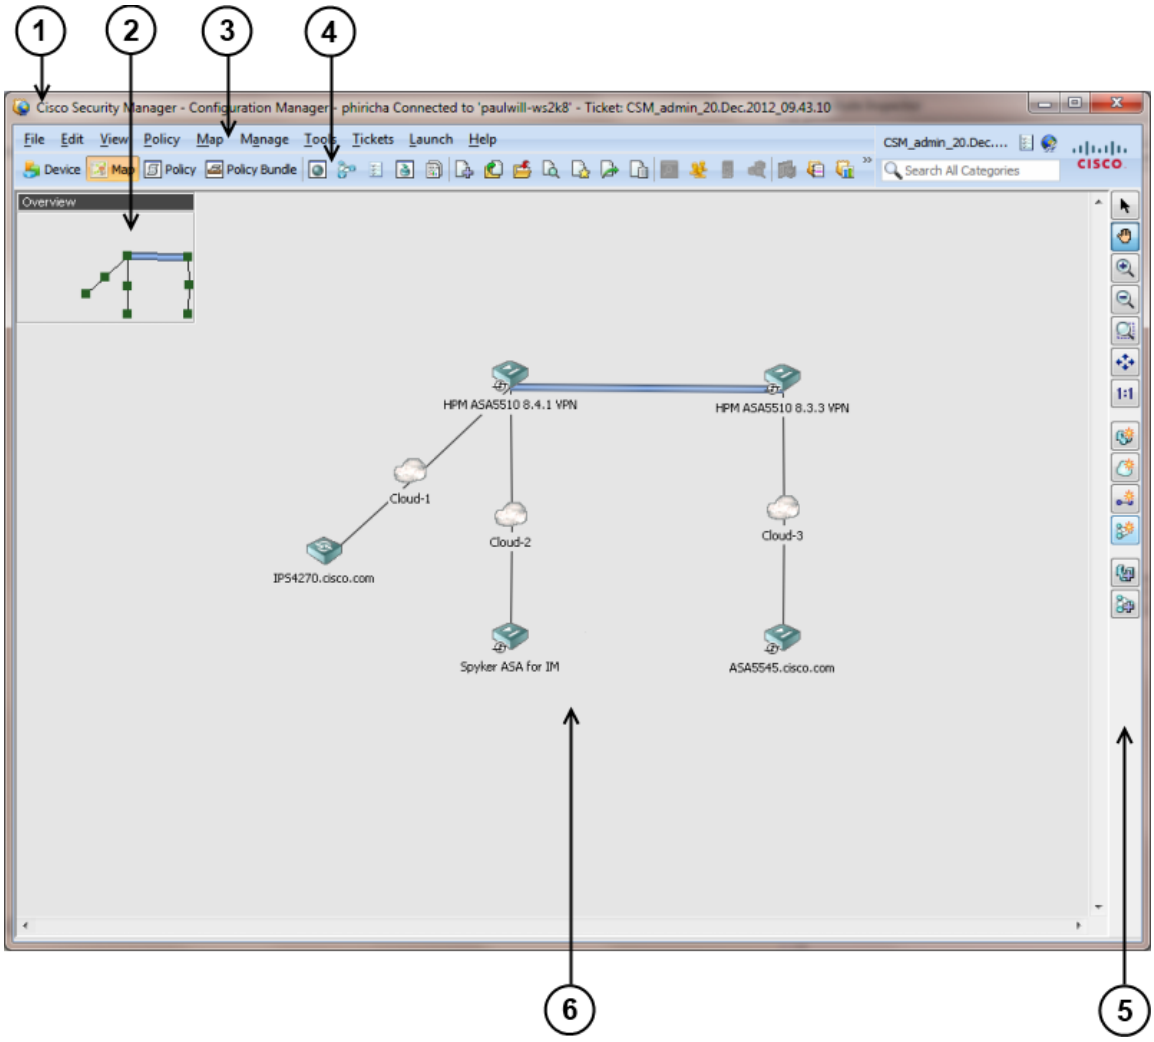

1	Title bar	**2**	Navigation window
3	Menu bar (see Map Menu (Configuration Manager), page 1-33)	**4**	Toolbar (see Toolbar Reference (Configuration Manager), page 1-39)
5	Map toolbar (see Map Toolbar, page 35-4)	**6**	Map

Task Flow for Configuring Security Policies

The basic user task flow for configuring security policies on devices involves adding devices to the Security Manager inventory, defining the policies, and then deploying them to the devices. You perform these tasks in Configuration Manager. The following briefly describes the steps in a typical user task flow:

Step 1 **Prepare devices for management.**

Before you can add a device to the Security Manager device inventory and manage it, you must configure some minimal settings on the device to enable Security Manager to contact it. For more information, see Chapter 2, "Preparing Devices for Management".

Step 2 **Add devices to the Security Manager device inventory.**

To manage a device with Security Manager, you must first add it to the Security Manager inventory. Security Manager provides multiple methods to add devices: from the network (live devices), from an inventory file exported from another Security Manager server or CiscoWorks Common Services Device Credential Repository (DCR), or in Cisco Security Monitoring, Analysis and Response System (CS-MARS) format, or from a device configuration file. You can also add a device that does not yet exist in the network but which will be deployed in the future, by creating it in Security Manager.

When you add a device, you can also discover its interfaces and certain policies that were already configured on the device. Discovery brings the information into the Security Manager database for continued management with Security Manager in the future.

For more information, see Chapter 3, "Managing the Device Inventory".

Step 3 **Define security policies.**

After you have added your devices, you can define the security policies you require. You can use Device view to define policies on specific devices. You can use Policy view to create and manage reusable policies that can be shared by any number of devices. When you make a change to a shared policy, the change is applied to all devices to which that policy is assigned.

To simplify and speed up policy definition, you can use policy objects, which are named, reusable representations of specific values. You can define an object once and then reference it in multiple policies instead of having to define the values individually in each policy.

Note If you are using Workflow mode, you must create an activity before you start defining policies. For more information, see Workflow and Activities Overview, page 1-20.

For more information, see these topics:

- Chapter 5, "Managing Policies"
- Chapter 6, "Managing Policy Objects"

Step 4 **Submit and deploy your policy definitions.**

Policy definition is done within your private view. Your definitions are not committed to the database and cannot be seen by other Security Manager users until you submit them. When you submit your policy definitions, the system validates their integrity. Errors or warnings are displayed to inform you of any problems that need to be addressed before the policies can be deployed to the devices.

Security Manager generates CLI commands according to your policy definitions and enables you to quickly and easily deploy them to your devices. You can deploy directly to live devices in the network (including dynamically addressed devices) through a secure connection, or to files that can be transferred to your devices at any time.

In non-Workflow mode, submitting and deploying your changes can be done in a single action. In Workflow mode, you first submit your activity and then you create a deployment job to deploy your changes.

For more information, see Chapter 8, "Managing Deployment".

Policy and Policy Object Overview

A **policy** is a set of rules or parameters that define a particular aspect of network configuration. In Configuration Manager, you define policies that specify the security functionality you want on your devices. Security Manager translates your policies into CLI commands that can be deployed to the relevant devices.

Security Manager enables you to configure local policies and shared policies.

- **Local policies** are confined to the device on which they are configured; they are automatically assigned (applied) to the device when you configure them. Unconfigured policies (those whose default settings you do not change) are not considered to be assigned or configured. To remove a policy, you unassign it.
- **Shared policies** are named, reusable policies that can be assigned to multiple devices at once. Any changes you make to a shared policy are reflected on all devices to which that policy is assigned, so you do not have to make the change on each device.

When you add a device to the inventory, you can discover the existing policies configured on the device. Security Manager translates your device configuration into Security Manager policies, populates the relevant local policies, and assigns them to the device. **Policy discovery** ensures that you do not need to recreate your existing configurations in Security Manager terms. You can also rediscover policies on devices after you add them to the inventory if you change their configuration through the CLI.

When you create policies, you often have the option to use **policy objects**, which are reusable definitions of related sets of values. (Sometimes, you are required to use policy objects.) For example, you can define a network object called MyNetwork that contains a set of IP addresses in your network. Whenever you configure a policy requiring these addresses, you can simply refer to the MyNetwork network object rather than manually entering the addresses each time. Furthermore, you can make changes to policy objects in a central location and these changes will be reflected in all the policies that reference those objects.

For more detailed information, see Understanding Policies, page 5-1 and Chapter 6, "Managing Policy Objects".

Workflow and Activities Overview

To provide flexible, secure policy management while allowing your organization to implement change control processes, Security Manager provides three closely-related features in Configuration Manager:

- **Workflow/Non-Workflow modes**—Configuration Manager provides two modes of operation that scale to different organizational working environments: Workflow mode and non-Workflow mode (the default).
  - **Workflow Mode**—Workflow mode is for organizations that have division of responsibility between users who define security policies and those who administer security policies. It imposes a formal change-tracking and management system by requiring all policy configuration to be done within the context of an explicitly-created activity. A user can create multiple

activities so that a single activity contains only logically-related policy changes. You can configure Workflow mode to require a separate approver, so that configuration changes cannot be made without oversight. After approval, the user defines a separate deployment job to push the policy changes to the devices. For more information, see Working in Workflow Mode, page 1-21.

- **Non-Workflow Mode**—In non-Workflow mode, you do not explicitly create activities. When you log in, Configuration Manager creates an activity for you or opens the one you were previously using if it was not submitted. You can define and save your policies, and then submit and deploy them in one step. For more information, see Working in Non-Workflow Mode, page 1-22.

For information on selecting a mode, see Changing Workflow Modes, page 1-28.

- **Activities or Configuration Sessions**—An activity (in non-Workflow mode, a configuration session), is essentially a private view of the Security Manager database. In Configuration Manager, you use activities to control changes made to policies and policy assignments. Adding devices to the inventory does not involve an activity, however, unless you discover policies that define security contexts (on multi-context firewall devices) or virtual sensors (on IPS devices). Isolating policy changes in activities helps prevent "work in progress" from accidentally making it into active device configurations. For more information about activities and configuration sessions, see Understanding Activities, page 4-1 and Working with Activities/Tickets, page 4-7.
- **Ticket Management**—Ticket management allows you to associate a Ticket ID with policy configuration changes made in Security Manager. Ticket management works in coordination with activities or configuration sessions depending on whether you have workflow mode enabled or not. If workflow mode is enabled, you can also enable ticket management so that a Ticket ID can optionally be associated with a specific activity. If workflow mode is not enabled, using ticket management makes it so that all changes must be done as part of a ticket and the ticket must be submitted before those changes can be deployed. In this respect, ticket management with workflow disabled is very similar to how activities function when workflow is enabled; however, no approval of submitted tickets is required.

For a comparison of the various modes of operation, see Comparing Workflow Modes, page 1-22.

Working in Workflow Mode

Workflow mode is an advanced mode of operation that imposes a formal change-tracking and change-management system. Workflow mode is suitable for organizations in which there is division of responsibility among security and network operators for defining policies and deploying those policies to devices. For example, a security operator might be responsible for defining security policies on devices, another security operator might be responsible for approving the policy definitions, and a network operator might be responsible for deploying the resulting configurations to a device. This separation of responsibility helps maintain the integrity of deployed device configurations.

You can use Workflow mode with or without an approver. When using Workflow mode with an approver, device management and policy configuration changes performed by one user are reviewed and approved by another user before being deployed to the relevant devices. When using Workflow mode without an approver, device and policy configuration changes can be created and approved by a single user, thus simplifying the change process.

Note Workflow mode works in the same manner whether Ticket Management is enabled or not. Enabling Ticket Management in Workflow mode simply enables the Ticket field for use with Activities. Entering a ticket ID is not required, but if one is used, the Ticket field can be configured to link to an external change management system. For more information, see Ticket Management.

For information about enabling or disabling Workflow mode or enabling or disabling Ticket Management, see Changing Workflow Modes, page 1-28.

In Workflow mode:

- A user must create an activity before defining or changing policy configurations in Configuration Manager. The activity is essentially a proposal to make configuration changes. The changes made within the activity are applied only after the activity is approved by a user with the appropriate permissions. An activity can either be submitted to another user for review and approval, or it can be approved by the current user. For detailed information about the process of creating, submitting, and approving activities, see Chapter 4, "Managing Activities".
- After the activity is approved, the configuration changes need to be deployed to the relevant devices. To do this, a user must create a *deployment job*. A deployment job defines the devices to which configurations will be deployed, and the deployment method to be used. A deployment job can either be submitted to another user for review and approval, or it can be approved by the current user. Deployment preferences can be configured with or without job approval. For more information, see Chapter 8, "Managing Deployment"

Working in Non-Workflow Mode

Some organizations have no division of responsibility between users when defining and administering their VPN and firewall policies. These organizations can work in non-Workflow mode. When using non-Workflow mode, you do not explicitly create activities. When you log in, Configuration Manager creates an activity for you, also called a configuration session, or opens the activity you were using when previously logged in (the configuration session is automatically closed when you log out of Security Manager). This activity is transparent to the user and does not need to be managed in any way. When you submit your configuration changes to the database, this is equivalent to submitting and approving the activity in Workflow mode. In addition, when you submit and deploy configuration changes, Security Manager creates a deployment job for you as well. Like activities, deployment jobs are transparent and do not need to be managed.

When using non-Workflow mode, multiple users with the same username and password cannot be logged into Security Manager at the same time. If another user logs in with the same username and password while you are working, your session will be terminated and you will have to log in again.

Ticket Management in Non-Workflow Mode

If your organization uses a change management system, Security Manager can associate the changes made to configurations with a ticket ID. Before making any configuration changes, you must open a ticket and the ticket must be submitted before the changes associated with that ticket are available to be deployed. Tickets can be opened and closed as needed, and you can discard a ticket if the changes associated with that ticket are no longer desired. Entering a ticket ID is not required, but if one is used, the Ticket field can be configured to link to an external change management system. For more information, see Ticket Management.

Non-Workflow mode with Ticket Management enabled is the default mode for Security Manager. For information about enabling or disabling Workflow mode or enabling or disabling Ticket Management, see Changing Workflow Modes, page 1-28.

Comparing Workflow Modes

The following table highlights the differences between the workflow modes.

Note Workflow mode works in the same manner whether Ticket Management is enabled or not. Enabling Ticket Management in Workflow mode simply enables the Ticket field for use with Activities. Entering a ticket ID is not required, but if one is used, the Ticket field can be configured to link to an external change management system. For more information, see Ticket Management.

Table 1-1 Comparison Between Workflow Mode and Non-Workflow Mode in Configuration Manager

Question	Non-Workflow Mode with Ticket Management Enabled	Non-Workflow Mode with Ticket Management Disabled	Workflow Mode
What is the default mode for Security Manager?	Default	Not Default	Not default
How do I know which mode is currently selected?	Select **Tools > Security Manager Administration > Workflow**. If the **Enable Workflow** check box is selected, you are in Workflow mode. Select **Tools > Security Manager Administration > Ticket Management**. If the **Enable Ticketing** check box is selected, ticket management is enabled.		
Must I explicitly create activities to make configuration changes?	You must explicitly create a Ticket before you can make configuration changes. Configuration Manager automatically creates an activity that is associated with that ticket.	No. Configuration Manager automatically creates an activity when you log in, or opens the previous session if you did not submit it before logging out.	Yes.
Must I explicitly create deployment jobs to deploy configurations to devices?	No. Configuration Manager creates a deployment job for you when you deploy configuration changes.	No. Configuration Manager creates a deployment job for you when you deploy configuration changes.	Yes.
How do I deploy my configuration changes to the devices?	Do one of the following: • Select **File > Deploy**. • Select **Manage > Deployments** and click **Deploy** on the Deployment Jobs tab.	Do one of the following: • Click the **Submit and Deploy Changes** button in the Main toolbar. • Select **File > Submit and Deploy**. • Select **Manage > Deployments** and click **Deploy** on the Deployment Jobs tab.	Select **Manage > Deployments** and create a deployment job.

Table 1-1 Comparison Between Workflow Mode and Non-Workflow Mode in Configuration Manager (continued)

Question	Non-Workflow Mode with Ticket Management Enabled	Non-Workflow Mode with Ticket Management Disabled	Workflow Mode
At what stage are the CLI commands for my configuration changes generated?	When initiating deployment.	When initiating deployment.	When creating a deployment job.
How do I delete my current changes?	Select **Tickets > Discard Ticket** to discard the currently-open ticket, or select the ticket in the Ticket Manager and click **Discard**. If you have already started deploying devices, abort the deployment by selecting the job in the Deployment Manager and clicking **Abort**.	Select **File > Discard**. If you have already started deploying devices, abort the deployment by selecting the job in the Deployment Manager and clicking **Abort**.	Select **Activities > Discard Activity** to discard the currently-open activity, or select the activity in the Activity Manager and click **Discard**. If you already created a deployment job, select the job in the Deployment Manager and click **Discard**. If the job has already been deployed, you can abort the job by selecting **Abort**.
Can multiple users log into Security Manager at the same time?	Yes. Each user can open a different ticket and make configuration changes. A single user can log in multiple times, but the user must open separate tickets.	Yes, but only if each one has a different username. If a user with the same username logs into Security Manager, the first user is automatically logged out.	Yes. Each user can open a different activity and make configuration changes. A single user can log in multiple times, but the user must open separate activities.
What if another user is configuring the devices I want to configure?	You will receive a message indicating that the devices are locked. See Activities and Locking, page 4-3.		

Using the JumpStart to Learn About Security Manager

The JumpStart is an introduction to Security Manager. It describes and illustrates the major concepts of using the product. Use the jumpstart to explore Security Manager features and capabilities.

The JumpStart opens automatically when you first launch Security Manager. To get to the JumpStart while you are working with Security Manager, select **Help > JumpStart** from the main menu in Configuration Manager.

The JumpStart contains the following navigation features:

- A table of contents, which is always visible in the upper right corner. Click an entry to open its page.

- Links in the page enable you to drill down to more detailed information in the JumpStart or to relevant information in the online help.

Completing the Initial Security Manager Configuration

After you install Security Manager, there are several configuration steps you might want to perform to complete the installation. Although most of the features you initially configure have default settings, you should familiarize yourself with the features and decide if the default settings are the best settings for your organization.

The following list explains the features you might want to initially configure, with pointers to topics that provide more detailed information where appropriate. You can configure these features in any order, or delay configuring those that you do not yet need to use.

- Configure an SMTP server and default e-mail addresses. Security Manager can send e-mail notifications for several actions that occur in the system. For example, you can get an e-mail when your deployment job finishes reconfiguring network devices. For e-mail notifications to work, you must configure an SMTP server.

  For information on configuring an SMTP server and setting the default e-mail addresses, see Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27

- Create user accounts. Users must log into Security Manager to use the product. However, if a user logs in with an account another user is already using, the first user is automatically disconnected. Thus, each user should have a unique account. You can create accounts local to the Security Manager server, or you can use your ACS system to manage user authentication. For more information, see the *Installation Guide for Cisco Security Manager*

- Configure default deployment settings. When users deploy configurations to devices, they can select how the configurations should be deployed and how Security Manager should handle anomalies. However, you can select system-default settings that make it easier for users to follow your organization's recommendations. To set deployment defaults, in Configuration Manager, select **Tools > Security Manager Administration**, and then select **Deployment** from the table of contents to open the Deployment settings page (see Deployment Page, page 11-13).

  The following deployment settings are of particular interest:

  - Default Deployment Method—Whether configuration deployments should be written directly to the device or to a transport server, or if configuration files should be written to a specified directory on the Security Manager server. The default is to deploy configurations directly to the device or transport server, if one is configured for the device. However, if you have your own methods for deploying configuration files, you might want to select File as the default deployment method. For more information on deployment methods, see Understanding Deployment Methods, page 8-8

  - When Out-of-Band Changes Detected—How to respond when Security Manager detects that configuration changes were made on the device through the CLI rather than through Security Manager. The default is to issue a warning and proceed with the deployment, overwriting the changes that were made through the CLI. However, you can change this behavior to simply skip the check for changes (which means Security Manager overwrites the changes but does not warn you), or to cancel the deployment, thus leaving the device in its current state. For more information about handling out-of-bound changes, see Understanding How Out-of-Band Changes are Handled, page 8-12.

  - Allow Download on Error—Whether to allow deployment to continue if minor configuration errors are found. The default is to not allow deployment when minor errors are found.

- Select a workflow mode. The default mode is non-Workflow mode with Ticket Management enabled. In non-Workflow mode, users have more freedom to create and deploy configurations. However, if your organization requires a more transaction-oriented approach to network management, where separate individuals perform policy creation, approval, and deployment, you can enable Workflow mode to enforce your procedures. If you are using Workflow mode, ensure that you configure user permissions appropriately when you define user accounts to enforce your required division of labor. For information on the types of workflow you can use, see Workflow and Activities Overview, page 1-20. For information on how to change workflow modes, see Changing Workflow Modes, page 1-28.

Tip You can disable Ticket Management in non-Workflow mode to make most activity management tasks automatic.

- Configure default device communication settings. Security Manager uses the most commonly used methods for accessing devices based on the type of device. For example, Security Manager uses SSH by default when contacting Catalyst switches. If the default protocols work for the majority of your devices, you do not need to change them. For devices that should use a non-default protocol, you can change the protocol in the device properties for the specific devices. However, if you typically use a protocol that is not the Security Manager default (for example, if you use a token management server (TMS) for your routers), you should change the default setting. To change the default communication settings, in Configuration Manager, select **Tools > Security Manager Administration**, and select **Device Communication** from the table of contents. In the Device Connection Settings group, select the most appropriate protocols for each type of device. You can also change the default connection time out and retry settings. For more information about device communication settings, see Device Communication Page, page 11-21
- Select the types of router and firewall policies you will manage with Security Manager. When you manage IPS devices in Security Manager, you automatically manage the entire configuration. However, with routers and firewall devices (ASA, PIX, and FWSM), you can select which types of policies are managed by Security Manager. You can manage other parts of the device configuration using other tools (including the devices's CLI). By default, all security-related policies are managed. To change which policies are managed, in Configuration Manager, select **Tools > Security Manager Administration > Policy Management**. For detailed information about changing these settings and what you should do before and after making the change, see Customizing Policy Management for Routers and Firewall Devices, page 5-11.
- Decide whether you want to use the Event Viewer to manage firewall and IPS events. You can configure the disk and location for collecting syslog events from devices, and the port number to use for syslog communication. If you do not want to use Security Manager for event management, you can turn off the feature, which is enabled by default. For more information on the configuration options, see Event Management Page, page 11-27.
- Configure Security Manager for communication with Cisco Security Monitoring, Analysis and Response System (CS-MARS). If you use CS-MARS for monitoring your network, you can identify the servers to Security Manager and then access CS-MARS event information from within Security Manager. For information on configuring this cross-communication, see Checklist for Integrating CS-MARS with Security Manager, page 72-36.

Configuring an SMTP Server and Default Addresses for E-Mail Notifications

Security Manager can send e-mail notifications for several types of events such as deployment job completion, activity approval, or ACL rule expiration. To enable e-mail notifications, you must configure an SMTP server that Security Manager can use for sending the e-mails. Then, you can configure e-mail addresses and notification settings on these settings pages (in Configuration Manager, select **Tools > Security Manager Administration** and select the page from the table of contents):

- Workflow page—For default e-mail addresses and notification settings for deployment jobs and activities. Users can override the defaults when managing deployment jobs and activities.
- Rules Expiration page—For default e-mail addresses and notification settings for ACL rule expiration. Rules expire only if you configure them with expiration dates.
- IPS Updates page—For the e-mail address that should be notified of IPS update availability.
- Server Security page—When you configure local user accounts (click **Local User Setup**), specify the user's e-mail address. This address is used as the default target for some notifications such as deployment job completion.
- Event Management page—When you configure an extended data storage location, you must specify at least one e-mail address. The email addresses receive notifications if problems arise with the use of the extended storage location. Also, if you are using the Syslog Relay Service, you can configure e-mail addresses that should be notified when the syslog relay service enters or exits CPU throttling.

Tip If you are using ACS for user authorization, you might have already configured an SMTP server and system administrator e-mail address in the ACS integration procedure as described in the *Installation Guide for Cisco Security Manager*. Security Manager sends a notification to this address if all ACS servers become unavailable.

Step 1 Access CiscoWorks Common Services on the Security Manager server:

- If you are currently using the Security Manager client, the easiest way to do this is to select **Tools > Security Manager Administration**, select **Server Security** from the table of contents, and click any button on that page (for example, **Local User Setup**).
- You can use your web browser to log into the home page on the Security Manager server (https://*servername*/CSCOnm/servlet/login/login.jsp) and click **Server Administration**.

Step 2 Click **Server > Admin** and select **System Preferences** from the table of contents.

Step 3 On the System Preferences page, enter the host name or IP address of an SMTP server that Security Manager can use. The SMTP server cannot require user authentication for sending e-mail messages.

Also, enter an e-mail address that CiscoWorks can use for sending e-mails. This does not have to be the same e-mail address that you configure for Security Manager to use when sending notifications. If you are using ACS for authorization, Security Manager sends an e-mail message to this address if all ACS servers become unavailable. This can alert you to a problem that needs immediate attention. The administrator might also receive e-mail messages from Common Services for non-ACS-related events.

Step 4 Click **Apply** to save your changes.

Changing Workflow Modes

You can change the workflow mode that Security Manager enforces if you have the appropriate administrator permissions. Changing the workflow mode has significant effects on users. Before making a change, be sure to understand the following:

- When you change the workflow mode, the change will take effect for all Security Manager users working from the same server.
- Before you can change from Workflow mode to non-Workflow mode, all activities in editable states (Edit, Edit Open, Submit, or Submit Open) must be approved or discarded, and all generated jobs must be deployed, rejected, discarded, or aborted so that the locks on the devices can be released. You do not have to do anything to jobs that are in the failed state.
- Before you can disable Ticket Management in non-Workflow mode, all tickets in editable states (Edit or Edit Open) must be submitted or discarded.
- If you change from Workflow mode to non-Workflow mode and then restore an earlier version of the database, Security Manager automatically changes to Workflow mode if the restored database has any activities in an editable state (Edit, Edit Open, Submit, or Submit Open). Approve or delete the editable activities, and then turn Workflow mode off again.
- When changing from non-Workflow mode to Workflow mode or enabling Ticket Management in non-Workflow mode, current configuration sessions are listed as activities/tickets in the Edit_Open state, and these activities/tickets must now be explicitly managed.
- When Ticket Management is enabled or disabled, any other users logged into Security Manager are logged out.

For an explanation of workflow modes, see Workflow and Activities Overview, page 1-20.

Step 1 In Configuration Manager, select **Tools > Security Manager Administration** and select **Workflow** from the table of contents to open the Workflow page (see Workflow Page, page 11-75).

Step 2 Configure the workflow mode settings in the Workflow Control group. If you select Enable Workflow (to use Workflow mode), you can also select these options:

- Require Activity Approval—To enforce explicit approval of activities before policy changes are committed to the database.
- Submitter can Approve Activity— Instead of separating submission and approval roles, a submitter can also approve his/her own activity, when enabled.
- Require Deployment Approval—To enforce explicit approval of deployment jobs before they can be run.
- Submitter can Approve Deployment Job—When enabled, submitter can approve deployment jobs submitted by him/her.

Step 3 Configure the e-mail notification settings. These are the default e-mail addresses for the e-mail sender (that is, Security Manager), the approvers, and another person or e-mail alias who should be notified when deployment jobs are complete.

You also have the options to include the job deployer when sending notifications of job status, and to require that e-mail notifications are sent for deployment job status changes.

Step 4 Click **Save** to save and apply changes.

Step 5 Select **Workflow** from the table of contents to open the Ticket Management page (see Ticket Management Page, page 11-72).

Step 6 Configure the Ticket Management settings. If you select Enable Ticketing, you can also select these options:

Note See Ticket Management Page, page 11-72 for detailed information on these fields.

- Ticket System URL—To provide linking between a Ticket ID and an external ticket management system.
- Ticket History—Specify how long to keep information related to tickets.

Step 7 Click **Save** to save and apply changes.

Understanding Basic Security Manager Interface Features

The following topics provide information about some basic interface features such as descriptions of the menu commands, toolbar buttons, and how to use common user interface elements. Many of the features described are used only in Configuration Manager.

- Menu Bar Reference for Configuration Manager, page 1-29
- Toolbar Reference (Configuration Manager), page 1-39
- Using Selectors, page 1-45
- Using Wizards, page 1-47
- Using Tables, page 1-48
- Using Text Fields, page 1-49
- Selecting or Specifying a File or Directory in Security Manager, page 1-50
- Troubleshooting User Interface Problems, page 1-51

Menu Bar Reference for Configuration Manager

The menu bar in Configuration Manager contains menus with commands for using Security Manager. Commands may become unavailable depending on the task you are performing.

The menus in the menu bar are described in the following topics:

- File Menu (Configuration Manager), page 1-30
- Edit Menu (Configuration Manager), page 1-31
- View Menu (Configuration Manager), page 1-31
- Policy Menu (Configuration Manager), page 1-32
- Map Menu (Configuration Manager), page 1-33
- Manage Menu (Configuration Manager), page 1-34
- Tools Menu (Configuration Manager), page 1-34
- Launch Menu (Configuration Manager), page 1-37
- Activities Menu (Configuration Manager), page 1-36
- Tickets Menu (Configuration Manager), page 1-36

- Help Menu (Configuration Manager), page 1-38

File Menu (Configuration Manager)

The following table describes the commands on the File menu in Configuration Manager. The menu items differ depending on the workflow mode.

Table 1-2 File Menu (Configuration Manager)

Command	Description
New Device	Initiates the wizard to add a new device. See Adding Devices to the Device Inventory, page 3-6.
Clone Device	Creates a device by duplicating an existing device. See Cloning a Device, page 3-56
Delete Device	Deletes a device. See Deleting Devices from the Security Manager Inventory, page 3-58.
Save	Saves any changes made on the active page, but does not submit them to the Security Manager database.
Import	Import policies and devices exported from another Security Manager server. See Importing Policies or Devices, page 10-13.
Export	Export policies or devices so that they can be imported into another Security Manager server. A device export can include policy information, or it can be a simple CSV file that you can import into CiscoWorks Common Services Device Credential Repository (DCR) or Cisco Security Monitoring, Analysis and Response System (CS-MARS). See Exporting the Device Inventory from the Security Manager Client, page 10-6 and Exporting Shared Policies, page 10-12.
View Changes (non-Workflow mode only)	Opens the Activity Change Report (in PDF format) for the current configuration session. To see changes for the current activity in Workflow mode, select **Activities > View Changes**.
Validate (non-Workflow mode only)	Validates the changes you have saved. See Validating an Activity/Ticket, page 4-18. To validate the current activity in Workflow mode, select **Activities > Validate Activity**.
Submit (non-Workflow mode only)	Submits all changes made since the last submission to the Security Manager database. To validate the current activity in Workflow mode, select **Activities > Submit Activity**.
Submit and Deploy (non-Workflow mode only)	Submits all changes made since the last submission to the Security Manager database and deploys all changes made since the last deployment. See Understanding Deployment, page 8-1. In Workflow mode, you must have your activity approved and then create a deployment job to deploy changes to devices.

Table 1-2 File Menu (Configuration Manager) (continued)

Command	Description
Deploy (non-Workflow mode only)	Deploys all changes made since the last deployment. See Understanding Deployment, page 8-1. In Workflow mode, you must have your activity approved and then create a deployment job to deploy changes to devices.
Discard (non-Workflow mode only)	Discards all configuration changes since the last submission. To validate the current activity in Workflow mode, select **Activities > Discard Activity**.
Edit Device Groups	Edits device groups. See Working with Device Groups, page 3-59.
New Device Group	Adds a device group. See Creating Device Groups, page 3-62.
Add Devices to Group	Adds a device to a group. See Adding Devices to or Removing Them From Device Groups, page 3-63.
Print	Prints the active page. Not all pages can be printed. If the Print command is not available, you cannot print the active page.
Exit	Exits Security Manager.

Edit Menu (Configuration Manager)

The following table describes the commands on the Edit menu in Configuration Manager. You can typically use these commands only when you are working with a table in a policy, and some work only for rules tables (see Using Rules Tables, page 12-8).

Table 1-3 Edit Menu (Configuration Manager)

Command	Description
Cut	Cuts the selected row in a rules table and saves it on the clipboard.
Copy	Copies the selected row in a rules table and saves it on the clipboard.
Paste	Pastes the rules table row from the clipboard to the into the rules table after the selected row.
Add Row	Adds a row into the active table.
Edit Row	Edits the selected table row.
Delete Row	Deletes the selected table row.
Move Row Up Move Row Down	Moves the selected row up or down in the rules table. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Global Search	Opens the Global Search window. For more information, see Using Global Search, page 1-42.

View Menu (Configuration Manager)

The View menu in Configuration Manager contains commands to navigate within the user interface or to alter the toolbar.

Table 1-4 View Menu

Menu Command	Description
Device View	Opens Device view. See Device View Overview, page 1-15.
Device Status View	Opens the Device Status View window. See Working with Device Status View, page 3-64.
Map View	Opens Map view. See Map View Overview, page 1-18.
Policy View	Opens Policy view. See Policy View Overview, page 1-16.
Policy Bundle View	Opens Policy Bundle view. See Managing Policy Bundles, page 5-57.
Customized Toolbar	Allows you to add or remove some optional buttons on the toolbar. For information on all the buttons that can appear on the toolbar, see Toolbar Reference (Configuration Manager), page 1-39.

Policy Menu (Configuration Manager)

The Policy menu in Configuration Manager contains commands for managing policies.

Table 1-5 Policy Menu (Configuration Manager)

Menu Command	Description
Share Policy	Saves the active local policy as a shared policy. See Sharing a Local Policy, page 5-41.
Unshare Policy	Saves the active shared policy as a local policy. See Unsharing a Policy, page 5-43.
Assign Shared Policy	Assigns shared policies to devices. See Assigning a Shared Policy to a Device or VPN Topology, page 5-44.
Unassign Policy	Unassigns the current policy from the selected device. See Unassigning a Policy, page 5-36.
Copy Policies Between Devices	Copies policies between devices. See Copying Policies Between Devices, page 5-33
Share Device Polices	Enables you to share local device policies. See Sharing a Local Policy, page 5-41.
Edit Policy Assignments	Edits assignment of shared policies to devices. See Modifying Policy Assignments in Policy View, page 5-54.
Clone Policy	Creates a copy of a policy with a new name. See Cloning (Copying) a Shared Policy, page 5-47.
Rename Policy	Renames a policy. See Renaming a Shared Policy, page 5-48.
Add Local Rules	Adds local rules to a shared policy on a device. You must select a rule-based shared policy to use this command.
Inherit Rules	Edits policy inheritance. See Inheriting or Uninheriting Rules, page 5-47
Discover Policies on Device	Discovers policies on a device. See Discovering Policies, page 5-12
Discover VPN Policies	Opens the Discover VPN Policies wizard. See Site-To-Site VPN Discovery, page 25-20.

Map Menu (Configuration Manager)

The Map menu in Configuration Manager contains commands for using the Map view. The commands in this menu are available only when the Map view is open. For more information, see Chapter 35, "Using Map View".

Table 1-6 Map Menu (Configuration Manager)

Menu Command	Description
New Map	Creates a map. See Creating New or Default Maps, page 35-9.
Open Map	Opens a saved map or the default map. See Opening Maps, page 35-10.
Show Devices On Map	Selects the managed devices to show on the active map. See Displaying Managed Devices on the Map, page 35-16.
Show VPNs On Map	Selects the VPNs to show on the active map. See Displaying Existing VPNs on the Map, page 35-21.
Add Map Object	Creates a map object on the open map. See Using Map Objects To Represent Network Topology, page 35-17.
Add Link	Creates a Layer 3 link on the open map. See Creating and Managing Layer 3 Links on the Map, page 35-19.
Find Map Node	Finds nodes on the open map. See Searching for Map Nodes, page 35-12.
Save Map	Saves the open map. See Saving Maps, page 35-10.
Save Map As	Saves the open map with a new name. See Saving Maps, page 35-10.
Zoom In	Zooms in on the map. See Panning, Centering, and Zooming Maps, page 35-11.
Zoom Out	Zooms out from the map. See Panning, Centering, and Zooming Maps, page 35-11.
Fit to Window	Zooms the open map to display the entire map. See Panning, Centering, and Zooming Maps, page 35-11.
Display Actual Size	Zooms the open map to display at actual size. See Panning, Centering, and Zooming Maps, page 35-11.
Refresh Map	Refreshes the open map with updated network data. See Creating New or Default Maps, page 35-9.
Export Map	Exports the open map to a file. See Exporting Maps, page 35-11.
Delete Map	Deletes the map you select from a list. See Deleting Maps, page 35-10.
Map Properties	Displays or edits properties for the open map. See Setting the Map Background Properties, page 35-13.
Show/Hide Navigation Window	Displays or hides the navigation window on the open map. See Using the Navigation Window, page 35-4.
Undock/Dock Map View	Undocks the maps window, allowing you to use other features while keeping the map open. If the window is already undocked, the Dock Map View command reattaches the window to the primary Security Manager window. See Understanding the Map View Main Page, page 35-2.

Manage Menu (Configuration Manager)

The Manage menu in Configuration Manager contains commands that start tools that run in a window separate from the Security Manager main interface. This enables you to access features without closing the page from which you are currently working.

Table 1-7 Manage Menu (Configuration Manager)

Menu Command	Description
Policy Objects	Opens the Policy Object Manager, where you can view all available objects grouped according to object type; create, copy, edit, and delete objects; and generate usage reports, which describe how selected objects are being used by other Security Manager objects and policies. For information see Policy Object Manager, page 6-4.
Site-to-Site VPNs	Opens the Site-to-Site VPN Manager, where you can configure site-to-site VPNs. See Chapter 25, "Managing Site-to-Site VPNs: The Basics"
Activities (Workflow mode only)	Opens the Activity Manager, where you can create and manage activities. See Activity/Ticket Manager Window, page 4-10.
Deployments	Opens the Deployment Manager, where you can deploy configurations and manage deployment jobs. See Chapter 8, "Managing Deployment"
Configuration Archive	Stores archived device configuration versions and allows you to view, compare, and roll back from one configuration to another. See Configuration Archive Window, page 8-23.
Policy Discovery Status	Opens the Policy Discovery Status window, where you can see the status of policy discovery and device import. See Viewing Policy Discovery Task Status, page 5-22.
IPS	Manage IPS device certificates, which are required for device communications.
Audit Report	Generates an audit report according to parameters set in the audit report page. See Using the Audit Report Window, page 10-21.
Change Reports (non-Workflow mode only)	Allows you to generate a report of changes to devices, shared policies, and policy objects for a previous configuration session. See Viewing Change Reports, page 4-16. To view changes for the current configuration session, select **File > View Changes**.

Tools Menu (Configuration Manager)

The Tools menu in Configuration Manager contains commands that start tools that run in a window separate from the Security Manager main interface. This enables you to access features without closing the page from which you are currently working.

Table 1-8 Tools Menu (Configuration Manager)

Menu Command	Description
Device Properties	Opens the Device Properties window, which provides general information about the device, including credentials, the group the device is assigned to, and policy object overrides. For more information, see Understanding Device Properties, page 3-6.
Detect Out of Band Changes	Analyzes devices to determine if their configurations have changed since the last time Security Manager deployed configurations. You can use this information to ensure that you do not loose important configuration changes. See Detecting and Analyzing Out of Band Changes, page 8-45.
Packet Capture Wizard	Opens the Packet Capture wizard, where you can set up a packet capture on an ASA device.
Ping, TraceRoute and NSLookup	Opens the Ping, TraceRoute, and NSLookup tool, where you can use these troubleshooting commands. Ping and traceroute run on managed devices, whereas NSLookup runs on your client workstation. See Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools, page 72-26.
IP Intelligence	Opens the IP Intelligence tool, where you can access various pieces of information about an IPv4 address, such as the fully qualified domain name (FQDN), geographic location information, and WHOIS information. For more information on the IP Intelligence tool, see IP Intelligence, page 72-33. Before you can use any of the IP Intelligence features, you must enable and configure those features on the IP Intelligence Settings page (see IP Intelligence Settings Page, page 11-41).
Wall	Opens the Wall window, where you can send messages to all users who are logged in on the same Security Manager server. First, however, it must be enabled on the Wall Settings page. See Wall Settings Page, page 11-77.
Show Containment	Shows security contexts or service modules for a device. See Showing Device Containment, page 3-56.
Inventory Status	Shows device summary information for all devices. See Viewing Inventory Status, page 72-12.
Catalyst Summary Info	Shows high-level system information, including any service modules, ports, and VLANs that Security Manager has discovered on the selected Catalyst switch. See Viewing Catalyst Summary Information, page 68-2.
Apply IPS Update	Manually applies IPS image and signature updates. See Manually Applying IPS Updates, page 44-7.
Preview Configuration	Displays the proposed changes, last deployed configuration, or current running configuration for specific devices. See Previewing Configurations, page 8-44.
Backup	Backs up the Security Manager database using CiscoWorks Common Services. See Backing up and Restoring the Security Manager Database, page 10-24.

Table 1-8 Tools Menu (Configuration Manager) (continued)

Menu Command	Description
Security Manager Diagnostics	Gathers troubleshooting information to send to the Technical Assistance Center (TAC) if they request it. See Creating Diagnostics Files for the Cisco Technical Assistance Center, page 10-28. **Tip** Beginning with Version 4.7 of Cisco Security Manager, you can select "Light Diagnostics" instead of the existing "General Diagnostics."
Security Manager Administration	Configures system-wide settings that control the functioning of Security Manager. For information, see Chapter 11, "Configuring Security Manager Administrative Settings".

Activities Menu (Configuration Manager)

The Activities menu in Configuration Manager contains commands for managing activities. It appears only when Workflow mode is enabled. For more detailed information about these commands, see Accessing Activity Functions in Workflow Mode, page 4-8.

Table 1-9 Activities Menu (Configuration Manager)

Menu Command	Description
New Activity	Creates a new activity. See Creating an Activity/Ticket, page 4-14.
Open Activity	Opens an activity. See Opening an Activity/Ticket, page 4-15.
Close Activity	Closes the open activity. See Closing an Activity/Ticket, page 4-16.
View Changes	Opens the Activity Change Report (in PDF format). See Viewing Change Reports, page 4-16.
Validate Activity	Validates the open activity. See Validating an Activity/Ticket, page 4-18.
Submit Activity	Submits the open activity. See Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20.
Approve Activity	Approves the open activity. See Approving or Rejecting an Activity (Workflow Mode), page 4-21.
Reject Activity	Rejects the open activity. See Approving or Rejecting an Activity (Workflow Mode), page 4-21.
Discard Activity	Discards the open activity. See Discarding an Activity/Ticket, page 4-22.

Tickets Menu (Configuration Manager)

The Tickets menu in Configuration Manager contains commands for managing tickets. It appears only when Ticket Management is enabled in non-Workflow mode. For more detailed information about these commands, see Accessing Ticket Functions in Non-Workflow Mode, page 4-9.

Table 1-10 Tickets Menu (Configuration Manager)

Menu Command	Description
New Ticket	Creates a new ticket. See Creating an Activity/Ticket, page 4-14.
Open Ticket	Opens an ticket. See Opening an Activity/Ticket, page 4-15.
Close Ticket	Closes the open ticket. See Closing an Activity/Ticket, page 4-16.
View Changes	Opens the Ticket Change Report (in PDF format). See Viewing Change Reports, page 4-16.
Validate Ticket	Validates the open ticket. See Validating an Activity/Ticket, page 4-18.
Submit Ticket	Submits the open ticket. See Understanding Activity/Ticket States, page 4-4.
Discard Ticket	Discards the open ticket. See Discarding an Activity/Ticket, page 4-22.

Launch Menu (Configuration Manager)

The Launch menu contains commands that start other applications.

Table 1-11 Launch Menu (Configuration Manager)

Menu Command	Description
Device Manager	Starts device managers for all supported devices, such as PIX security appliances, Firewall Services Modules (FWSM), IPS sensors, IOS routers, and Adaptive Security Appliance (ASA) devices. Device managers provide several monitoring and diagnostic features that enable you to get information regarding the services running on the device and a snapshot of the overall health of the system. See Starting Device Managers, page 72-14.
Prime Security Manager	Launches the Cisco Prime Security Manager (PRSM) application, used to manage ASA CX devices. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.
FireSIGHT Management Center	Launches the FireSIGHT Management Center application, used to manage FirePOWER modules. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.
Dashboard	Opens the Dashboard, which is a configurable launch point for Security Manager that makes IPS and FW tasks more convenient for you. In addition to the original dashboard, you can create new, additional dashboards, and you can customize all dashboards. By using the dashboard, you can accomplish in one place many tasks that are found in several other areas of Security Manager, such as the IPS Health Monitor page, Report Manager, Health and Performance Monitor, and IP Intelligence Settings. For detailed information on the dashboard, see Dashboard Overview, page 72-1.

Table 1-11 Launch Menu (Configuration Manager) (continued)

Menu Command	Description
Event Viewer	Opens the Event Viewer, where you can view and analyze device events. See Chapter 69, "Viewing Events" for more information. If you have already logged into another Security Manager application, Event Viewer is opened using the same user account; you are not prompted to log in. To open Event Viewer using a different user account, open the application from the Windows Start menu or desktop icon.
Report Manager	Opens the Report Manager, where you can generate and analyze security and usage reports. See Chapter 70, "Managing Reports" for more information. If you have already logged into another Security Manager application, Report Manager is opened using the same user account; you are not prompted to log in. To open Report Manager using a different user account, open the application from the Windows Start menu or desktop icon.
Image Manager	Opens the Image Manager, where you can manage the images on ASA devices. See Chapter 73, "Using Image Manager" for more information. If you have already logged into another Security Manager application, Image Manager is opened using the same user account; you are not prompted to log in. To open Image Manager using a different user account, open the application from the Windows Start menu or desktop icon.
Health & Performance Monitor	Opens the Health & Performance Monitor (HPM), where you can view device status and traffic information across your network, and view and acknowledge device-specific alerts. See Chapter 71, "Health and Performance Monitoring" for more information. If you have already logged into another Security Manager application, HPM is opened using the same user account; you are not prompted to log in. To open HPM using a different user account, open the application from the Windows Start menu or desktop icon.

Help Menu (Configuration Manager)

The Help menu in Configuration Manager contains commands for accessing product documentation and training. For more information, see Accessing Online Help, page 1-52.

Table 1-12 Help Menu (Configuration Manager)

Menu Command	Description
Help Topics	Opens the online help system.
Help About This Page	Open online help for the active page.
JumpStart	Opens the JumpStart.
Security Manager Online	Opens the Security Manager web page on Cisco.com.

Table 1-12 Help Menu (Configuration Manager) (continued)

Menu Command	Description
About Configuration Manager	Displays information about Configuration Manager.

Toolbar Reference (Configuration Manager)

The main toolbar (see the illustration Figure 1-1) contains buttons that perform actions in Configuration Manager.

The buttons that appear on the main toolbar vary depending on whether Workflow/Ticket Management mode is enabled and how you have customized the toolbar. By selecting **View > Customized Toolbar**, you can select some of the buttons included in the toolbar. Many buttons are on the toolbar permanently; you cannot remove them.

The following table presents all buttons.

Table 1-13 Configuration Manager Toolbar

Button	Description
	Opens the Device view. For more information, see Understanding the Device View, page 3-1.
	Opens the Map view. For more information, see Chapter 35, “Using Map View”.
	Opens the Policy view. For more information, see Managing Shared Policies in Policy View, page 5-50.
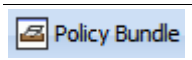	Opens the Policy Bundle view. For more information, see Managing Policy Bundles, page 5-57.
	Opens the Policy Object Manager. For more information, see Chapter 6, “Managing Policy Objects”.
	Opens the Site-to-Site VPN Manager. For more information, see Chapter 25, “Managing Site-to-Site VPNs: The Basics”.
	Opens the Deployment Manager. For more information, see Chapter 8, “Managing Deployment”.
	Opens the Audit Report. For more information, see Understanding Audit Reports, page 10-19.
	(Non-Workflow mode with Ticket Management disabled only.) Submits and deploys changes. For more information, see Chapter 8, “Managing Deployment”.
	Discovers configuration policies defined on the currently selected device. For more information, see Discovering Policies, page 5-12.

Table 1-13 Configuration Manager Toolbar (continued)

Button	Description
	Detects out-of-band changes, those made to the device outside of Security Manager, for the currently selected devices. For more information, see Detecting and Analyzing Out of Band Changes, page 8-45.
	Opens the IP Intelligence tool, where you can access various pieces of information about an IPv4 address, such as the fully qualified domain name (FQDN), geographic location information, and WHOIS information. For more information on the IP Intelligence tool, see IP Intelligence, page 72-33. Before you can use any of the IP Intelligence features, you must enable and configure those features on the IP Intelligence Settings page (see IP Intelligence Settings Page, page 11-41).
	Opens the Wall window, where you can send messages to all users who are logged in on the same Security Manager server. First, however, it must be enabled on the Wall Settings page. For more information, see Workflow Page, page 11-75.
	Shows high-level system information, including any service modules, ports, and VLANs that Security Manager has discovered on the selected Catalyst switch. For more information, see Viewing Catalyst Summary Information, page 68-2.
	Previews the configuration for the currently selected device. For more information, see Previewing Configurations, page 8-44.
	Configures system-wide settings that control the functioning of Security Manager. For information, see Chapter 11, “Configuring Security Manager Administrative Settings”.
	Opens the device manager for the currently selected device. For more information, see Starting Device Managers, page 72-14.
	Launches the Cisco Prime Security Manager (PRSM) application, used to manage ASA CX devices. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.
	Launches the FireSIGHT Management Center application, used to manage FirePOWER modules. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.
	Opens the Dashboard application. For more information, see Dashboard Overview, page 72-1.
	Opens the Event Viewer application. For more information, see Chapter 69, “Viewing Events”.
	Opens the Report Manager application. For more information, see Chapter 70, “Managing Reports”.
	Opens the Image Manager application. For more information, see Chapter 73, “Using Image Manager”.

Table 1-13 Configuration Manager Toolbar (continued)

Button	Description
	Opens the Health & Performance Monitor application. For more information, see Chapter 71, "Health and Performance Monitoring".
	Opens online help for the current page. For more information, see Accessing Online Help, page 1-52.
Note The following buttons are not available in non-Workflow mode when Ticket Management is disabled.	
	Opens the Activity Manager window in Workflow mode or the Ticket Manager window when Ticket Management is enabled in non-Workflow mode. You can use these windows to create and manage activities/tickets. For more information, see Activity/Ticket Manager Window, page 4-10. For more information on the activity buttons, and the conditions under which they are enabled, see Accessing Activity Functions in Workflow Mode, page 4-8. For more information on the ticket buttons, and the conditions under which they are enabled, see Accessing Ticket Functions in Non-Workflow Mode, page 4-9.
	Creates a new activity/ticket.
	Opens an activity/ticket.
	Saves all changes made while the activity/ticket was open and closes it.
	Evaluates all changes made in the activity/ticket and produces a Change Report in PDF format in a separate window. For more information, see Viewing Change Reports, page 4-16
	Validates the integrity of changed policies within the current activity/ticket.
	(Workflow mode with an approver only.) Submits the activity for approval when using Workflow mode with an activity approver. (Non-Workflow mode with Ticket Management enabled only.) Submits the ticket. Submitting the ticket saves the proposed changes to the database. Devices associated with the ticket are unlocked, meaning they can be included in policy definitions and changes in other tickets. You can submit a ticket when it is in the Edit or the Edit Open state.
	(Workflow mode only.) Approves the changes proposed in an activity.
	(Workflow mode only.) Rejects the changes proposed in an activity.
	Discards the selected activity/ticket.

Using Global Search

Security Manager provides a global search feature to make finding and working with information that you are interested in easier. The Global Search feature allows you to search for devices, policy objects, policies, and tickets that contain a particular search string. The scope of the search can be limited to just devices, policy objects, policies, or tickets.

Note Search is only performed using data that has been committed. Changes that have not yet been submitted to the database will not be included in search results.

Wildcard Matching

The search string supports the use of the following wildcard characters:

- **Asterisk (*)**—matches zero or more characters
- **Question Mark (?)**—matches a single character

Semantic Searching

If the search string that is entered is an IP address, Security Manager will perform a semantic search. For example, entering "192.168.0.0/16" in the search string will return items matching that subnet as well as any specific hosts or other subnets belonging to that subnet or to which that subnet belongs.

Global Search Scope

Global search is supported only within a set of policies and policy objects, not all. The supported policies and the policy objects are the most frequently used policies and objects in the customer deployments. The policies and policy objects supported are:

- Devices: All Devices
- Policy Objects:
  - AAA Server Groups
  - AAA Servers
  - Access Control Lists
  - As Path Policies
  - ASA Group Policies
  - BFD Template
  - Categories
  - Cisco Secure Desktop (Router)
  - Community List Policies
  - Credentials
  - DHCPv6 Pool
  - File Objects
  - FlexConfigs
  - Identity User Group
  - IKE Proposals
  - Interface Roles

  - IPSec Transform Sets
  - LDAP Attribute Maps
  - Networks/Hosts (IPv4 and IPv6)
  - PKI Enrollments
  - Policy List Policies
  - Port Forwarding List
  - Prefix List Policies
  - Route List Policies
  - Services
  - Single Sign On Servers
  - SLA Monitors
  - SSL VPN Bookmarks
  - SSL VPN Customizations
  - SSL VPN Gateways
  - SSL VPN Smart Tunnel Auto Signon Lists
  - SSL VPN Smart Tunnels
  - Text Objects
  - Time Ranges
  - Traffic Flows
  - User Groups
  - WINS Server Lists
- Policies:
  - AAA Rules
  - Access Rules
  - IPv6 Access Rules
  - Inspection Rules
  - Translation Rules
  - Web Filter Rules
  - Zone Based Firewall Rules
- Tickets
  - Configuration Manager
  - Image Manager

Performing a Global Search

To perform a global search, do one of the following:

- Select **Edit > Global Search** or press **Ctrl+F** to open the Global Search window. Select the scope for the search in the drop-down list to the left of the search field, enter your search string in the search field, and then click **Search**.

Note If you are currently viewing a rule table, pressing **Ctrl+F** will open the Find and Replace dialog box instead of the Global Search window. Use one of the other methods to access the Global Search feature instead of the Find and Replace feature.

- Using the search field in the upper-right corner of the Configuration Manager window, select the scope for the search by clicking on the Search icon, enter your search string in the search field, and then press **Enter**.

The Global Search window displays the results matching your search criteria. Select the desired data type from the Category selector tree to see results for that category.

Acting on Search Results

You can perform the following actions on the items returned from your search:

- **Export Data (All)**—Allows you to export the search results for the selected category in CSV format. Select the desired data type from the Category selector tree in the Global Search window to see results for that category, then click **Export** in the toolbar above the search results to export that table of data in CSV format.
- **Print (All)**—Allows you to print the search results for the selected category. Select the desired data type from the Category selector tree in the Global Search window to see results for that category, then click **Print** in the toolbar above the search results to print the table of data.
- **Device Properties (Devices)**—Allows you to view the device properties for devices returned in search results. Select the desired device group from the Category selector tree in the Global Search window to see results for that category. Select a device in the results table to highlight it, right-click the device, and then select **Device Properties**. The Device Properties dialog box for the selected device is displayed. For more information, see Viewing or Changing Device Properties, page 3-40.
- **Go To (Policies)**—Allows you to navigate to a policy from the search results. Select the desired policy type from the Category selector tree in the Global Search window to see results for that policy type. Select an item in the results table to highlight it, right-click the item, and then select **Go To**. The relevant policy for the selected item is displayed.
- **Filter (Policies)**—Allows you to filter the search results using the standard table filter. For more information, see Filtering Tables, page 1-48.
- **View (Policy Objects)**—Allows you to view the policy object details for an object in the search results. Select the desired policy object type from the Category selector tree in the Global Search window to see results for that object type. Select an object in the results table to highlight it, then click **View** in the toolbar above the search results (or right-click the object and select **View**). The relevant Edit dialog box for the selected policy object is displayed in read-only mode.
- **Edit (Policy Objects)**—Allows you to edit a policy object from the search results. Select the desired policy object type from the Category selector tree in the Global Search window to see results for that object type. Select an object in the results table to highlight it, then click **Edit** in the toolbar above the search results (or right-click the object and select **Edit**). The relevant Edit dialog box for the selected policy object is displayed.

Note If a ticket or activity is not currently open, you will be prompted to create one or open an existing one before you can edit the policy object.

- **Find Usage (Policy Objects)**—Allows you to find which policies, objects, VPNs, and devices are using an object in the search results. Select the desired policy object type from the Category selector tree in the Global Search window to see results for that object type. Select an object in the results

table to highlight it, then click **Find Usage** in the toolbar above the search results (or right-click the object and select **Find Usage**). The Object Usage dialog box for the selected policy object is displayed. For more information, see Generating Object Usage Reports, page 6-15.

- **Show Ticket (Tickets)**—Allows you to navigate to the Ticket Manager window for a ticket returned in the search results. Select the desired ticket group from the Category selector tree in the Global Search window to see results for that category. Click the Ticket column in the results table for the ticket you want to view. The Ticket Manager window is displayed with the selected ticket highlighted. For more information, see Activity/Ticket Manager Window, page 4-10.

Using Selectors

Selectors appear in several places in the user interface; for example, the Device selector in Device view (see Figure 1-1). These tree structures enable you to select items (like devices) on which to perform actions. Several types of items can appear in a selector, depending on the task you are performing.

Items in selectors are presented in a hierarchy of folders. You can browse for items in a selector by expanding and collapsing folders, which can contain other folders, items, or a combination of folders and items. To expand and collapse a folder, click the +/- next to it.

To select an item, click it. If it is possible to perform actions on multiple items (for example, in a device selector), you can use Ctrl+click to select each item, or Shift+click on the first and last item to select all items between them. Many selectors support auto select, that is, when you type a single letter, the next folder or item in the selector that begins with that letter is selected.

You can right-click an item to see commands that you can use with the item. Some commands on the right-click menus are unique and not repeated on the regular menus.

Many times a device selector appears in a dialog box divided into two panes, Available Devices and Selected Devices. In these dialog boxes, you must select the devices in the available devices list and click **>>** to move them to the selected list to actually select the devices. To deselect the devices, you select them in the selected devices list and click **<<**.

If a selector contains a large number of items, you can filter it to view a subset of those items. For more information, see Filtering Items in Selectors, page 1-45.

Filtering Items in Selectors

To view a subset of the items in a selector, you can create filters to display only those items that match the criteria you specify. You can have a maximum of 10 filters per user for each selector. After that, when you create another filter, that new filter replaces the oldest filter. There is no duplication check for filters that are created. You cannot delete filters manually.

A filter list appears above all selectors that can be filtered. From this list, you can do the following:

- Select a filter that you created previously.
- Select **None** to see the tree without any filters applied to it.
- Select **Create Filter** to create a filter.

Each filter can contain several filter rules. Each filter rule specifies a rule type, criteria, and values. You select whether items must match any or all filter rules before they can be displayed in the selector.

When you create a filter, the fields that you can filter on depend on the types of items displayed in the filter. However, the general procedure is the same for all selectors.

For information on filtering tables, see Filtering Tables, page 1-48.

Tip When you filter a selector, that filter might remain applied to the selector when you open another window that includes the selector. For example, when you apply a filter to the Device selector in Device view, that filter is applied to the selector if you open the New Device wizard. If you have problems finding an item in a selector, check the Filter field to see if a filter is being applied.

Step 1 Select **Create Filter** from the selector filter field to open the Create Filter dialog box.

Step 2 Select one of the radio buttons to determine the matching criteria. The choices are:

- Match Any of the Following—Creates an OR relationship among the filter criteria. Policies matching any of your criteria are included in the filter.
- Match All of the Following—Creates an AND relationship among the filter criteria. Only those policies matching all your criteria are included in the filter.

Step 3 Establish a filter rule by entering three criteria, as follows:

- From the first list, select the type to be filtered; for example, *Name*.
- From the next list, select the operating criteria for the filter; for example, *contains*.
- In the final field, enter or select a value on which to filter; for example *Cisco*.

Step 4 Click **Add**.

Tip If you make a mistake in forming the filter rule, select the rule and click **Remove** to delete it.

Step 5 Add any additional filter rules that you require. Click **OK** when you are finished.

The selector is filtered according to the new filter criteria, and the new filter is added to the filter list.

Create Filter Dialog Box

Use the Create Filter dialog box to filter and display a subset items in a selector or a table. Creating filters helps you find items more easily when viewing large lists.

For more information on filtering, see these topics:

- Filtering Items in Selectors, page 1-45
- Filtering Tables, page 1-48

Navigation Path

Do one of the following:

- Select **Create Filter** from the Filter field in a selector tree.
- Select **Advanced Filter** from the Filter field above a table.

Field Reference

Table 1-14 Create Filter Dialog Box

Element	Description
Match All of the Following	When you select this option an AND relationship is created among the filtering criteria you define. An item must satisfy every rule in the filter to be displayed in the list. For example, if you define the following criteria: • Name contains OSPF • Name contains West When you click OK, the filter is defined as: Name contains OSPF **and** Name contains West.
Match Any of the Following	When you select this option an OR relationship is created among the filtering criteria you define. An item must satisfy only one of the rules in the filter to be displayed in the list. For example, if you define the following criteria: • Name contains OSPF • Name contains RIP When you click OK, the filter is defined as: Name contains OSPF **or** Name contains RIP.
Filter Type (First field.)	The type of property on which you are filtering. For tables, this is the column heading. You might have only one option for filtering certain lists (for example, you might only be able to filter by the name of the item).
Filter Operator (Second field.)	The relationship between the filter type and the filter value. The available options depend on the selected type.
Filter Value (Third field.)	The value on which you want to filter. Depending on the selected type, you either enter a text string in this field, or you select a value from the list.
Filter Content Area Add button Remove button	The filter type, operator, and value that you have selected for each criterion. • To add a criterion, create it in the fields above this area and click **Add**. • To remove a criterion, select it and click **Remove**.

Using Wizards

Some tasks that you can perform with Security Manager are presented as wizards. A wizard is a series of dialog boxes (or steps) that enables you to perform a task. The current step number and the total number of steps in the wizard are displayed in the wizard title bar.

Wizards share the following buttons:

- **Back**—Returns to the previous dialog box. Enables you to review and modify settings that you defined in previous wizard steps.

- **Next**—Continues to the next dialog box. If this button is unavailable, you must define some required settings in the current dialog box before you can continue. Required settings are marked with an asterisk (*).
- **Finish**—Finishes the wizard, saving the settings you defined. You can finish the wizard whenever this button is available. If this button is not available, you must define more settings.
- **Cancel**—Closes the wizard without saving any settings.
- **Help**—Opens online help for the wizard.

Using Tables

Many policies in Security Manager use tables. A small number of policies use a specialized type of table called a rules table. Rules tables have extra features compared to standard tables; for more information, see Using Rules Tables, page 12-8.

Standard tables include these basic features:

- Table filter—You can filter the rows displayed to help you find items in a large table. For more information, see Filtering Tables, page 1-48.
- Table column headings—You can sort by column and move, show, and hide columns. For more information, see Table Columns and Column Heading Features, page 1-49.
- Table buttons—Use the buttons below the table to do the following:
  - Add Row button (+ icon)—Click this button to add an item to the table.
  - Edit Row button (pencil icon)—Select a row and click this button to edit its properties.
  - Delete Row button (trash can icon)—Select a row and click this button to delete it from the table.

Filtering Tables

You can filter the items in a table to view a subset that satisfies specific criteria. Filtering a table does not change the contents of the table, but allows you to focus on just those entries that currently interest you. This is helpful for tables that have hundreds of entries.

To filter a table, use the Filter fields above the table. With these controls, you can do the following:

- To do simple filtering, select the column name on which you want to filter, select the relationship you are looking for (such as “begins with”), enter the desired text string (or in some cases, select one of the pre-defined options), and click **Apply**.

  You can filter the results by selecting another criteria and clicking Apply. Your filters are added together, showing the results that satisfy all criteria. For example, you could first enter “Service begins with IP,” click Apply, then enter “Source contains 10.100.10.10,” and click Apply. The result would be a table that shows all rows where the service is IP AND the source includes 10.100.10.10 (it might include other IP addresses as well).
- To do advanced filtering, select **Advanced Filter** from the left most menu (the one that contains the column headings). This opens the Create Filter dialog box. Using this dialog box, you can create multiple filter criteria just as you can with the regular filter controls. However, you also have the option to create a list of disjoined, OR’ed criteria, by selecting **Match Any of the Following**, where you can say “show me all rows that have IP for service or 10.100.10.10 for source address.”
  - To add criteria, enter the criteria and click **Add**.

  - To remove criteria, select the undesired criteria and click **Remove**.

  If you filter a table using the simple method, you can select Advanced Filter to alter your existing filter, adding or removing criteria as desired. The dialog box is filled with whatever filter criteria are currently applied to the table.

- The current filter is shown next to the Filter label in the filter control area. You can click **Clear** to remove the filter and show all rows.
- Any filter you apply is kept in the left most menu below the Advanced Filter entry. You can apply the filter by selecting it from the list. However, this list can have at most 10 entries. When you create your eleventh filter, your oldest filter is removed from the list. If you select a filter and add criteria, you are modifying that filter rather than creating a new one. You cannot delete the listed filters.

Tip Your filter is maintained for a given type of table even if you select another device or log out and subsequently log back in. For example, if you filter the Access Rules table for one device, it will be filtered the same way for other devices. When you clear the filter, it is cleared for the same type of table for all devices. Your filters do not affect what any other user sees.

Table Columns and Column Heading Features

Tables contain columns, each of which has a column heading in the heading row. These columns and their headings include the following features:

- Show/hide Columns—Right-click the table heading row to open the context menu and then select Show Columns. This menu enables you to select which columns appear. Showing or hiding columns does not affect the content of items defined in the table; it affects only your view.

  By default, the tables for some policies do not display all available columns.
- Show Details/Show Summary—Right-click the table heading row to open the context menu and then select either **Show Details** or **Show Summary**. This toggling menu enables you to select whether to view detailed or summarized information in the table.
- Move columns—Click and drag a column heading to move the column to a new position.
- Resize columns—Click a column heading divider (when the cursor turns into an arrow) and drag it to resize the column.
- Sort by column headings—Click a column heading to sort the table by that column's contents. Click the same column heading again to reverse the sort order. The sorted column has an arrow next to its heading.

Using Text Fields

Text fields can be single- or multiple-line, depending on the purpose of the field. Text fields that can contain multiple text lines include several features to make them easier to use. The following topics describe limitations and features of text fields:

- Understanding ASCII Limitations for Text, page 1-50
- Finding Text in Text Boxes, page 1-50
- Navigating Within Text Boxes, page 1-50

Understanding ASCII Limitations for Text

Devices typically restrict text to ASCII characters. If you include non-ASCII characters in Security Manager text fields that are used to generate commands in a device configuration file, the presence of those characters can prevent the configuration file from loading on the device. For example, a non-ASCII character in an interface description for an FWSM can prevent the device from loading the startup configuration when you restart the device.

The only places where you can include non-ASCII, non-English languages in device configurations is in the SSL VPN Bookmarks and SSL VPN Customization policy objects, which are used in configuring browser-based clientless SSL VPNs on ASA devices. For information on how you can support local languages for these objects, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

Finding Text in Text Boxes

Use the Find dialog box to find text within a multiple line text field.

Step 1 Click in a multiple line text field.

Step 2 Press **Ctrl+F**. The Find dialog box opens.

Step 3 Enter text to search for in the Find what field.

Step 4 To specify the direction of the search, select either **Up** or **Down** in the Direction field.

Step 5 To match the case of the text you entered, select the **Match Case** check box.

Step 6 Click **Find**. The next occurrence of your search text is highlighted in the text field.

Navigating Within Text Boxes

Use the Goto line dialog box to navigate to a specific line in a multiple line text field.

Step 1 Click in a multiple line text field.

Step 2 Press **Ctrl+G**. The Goto line dialog box opens.

Step 3 Enter a line number in the Line number field.

Step 4 Click **OK**. The text field scrolls to the line number you entered.

Selecting or Specifying a File or Directory in Security Manager

Cisco Security Manager uses a standard file system browser to let you select a directory or file or to specify a file.

You will be able to choose between client and server file systems when performing the following file operations:

- Installing Security Manager license files
- Importing/exporting device inventory files
- Importing/exporting shared policies

- Creating the following file objects:
  - Cisco Secure Desktop Package
  - Plug-In—For browser plug-in files.
  - AnyConnect Profile
  - AnyConnect Image
  - Hostscan Image

For all other file operations, you can create or select files only on the Security Manager server--you cannot use a drive mounted on the server, and you cannot use your client system.

Tip You can control whether file operations are allowed on the Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10.

Typically, to create or select a file, you click a **Browse** button to open a dialog box that has a title related to the action you are performing (for example, Choose Files when selecting configuration files). The Browse button appears on various dialog boxes throughout the product.

In the dialog box, use the folder tree on the left to navigate to the folder you want:

- If client-side file browsing is enabled and you are performing a function that supports client-side browsing (see above), select the tab that corresponds to the system you want to import from or export to.
- If you are selecting a file, find it in the folder tree and select it in the right pane. If the action you are taking allows you to select multiple files, use Ctrl+click to select files individually, or Shift+click to select a range of files. You might also need to select a file type to view only those files that apply to your action.
- If you are specifying (creating) a file, navigate to the folder in which you want to create the file, enter a file name, and select the appropriate file type.

Note The path and file name are restricted to characters in the English alphabet. Japanese characters are not supported. When selecting files on a Windows Japanese OS system, the usual file separator character \ is supported, although you should be aware that it might appear as the Yen symbol (U+00A5).

Troubleshooting User Interface Problems

The following tips might help you resolve general user interface problems that you might encounter:

- **Interface appears to freeze**—Occasionally, when you go from a Security Manager dialog box to some other application (for example, to check your e-mail), when you come back to Security Manager, nothing you click on responds. It appears the interface is frozen.

  This might be caused by an open dialog box that is covered by another Security Manager window. Until you close the dialog box, you will not be able to use any other window in the application. To find the hidden dialog box, press Alt+Tab, which opens a Windows panel that has icons for all currently open windows. Keep holding Alt, then press Tab repeatedly to cycle through the icons until you find the right one (the icon might be a generic Java icon rather than the Security Manager icon). You can also use your mouse to click the desired icon rather than using Tab to cycle through them.

- **Text and list elements missing, Java errors when clicking buttons**—If you change your Windows color scheme while running the Security Manager client, you must close and then restart the client. Otherwise, the behavior of the client can be unpredictable.

  If you are experiencing these problems and you did not change the color scheme, try closing and restarting the application.

- **Dialog Box is too big for the screen**—The minimum screen resolution for the Security Manager client is actually bigger than the best screen resolution available on many laptops (for screen resolution requirements, see the client system requirements in the *Installation Guide for Cisco Security Manager*). Because some dialog boxes are quite large, if you run the client on a laptop, you might find the occasional dialog box that is too big to fit on your screen.

  Usually, you can reposition the dialog box to get access to the OK, Cancel, and Help buttons. However, if you cannot get those buttons on the screen, you can use the following techniques to perform the same actions:

  – **OK**—Put your cursor in a field near the bottom of the dialog box, then press Tab to move from field to field. Typically, the first off-screen field is the OK button. When the cursor highlight moves off screen, press Enter.

    You can also put the cursor in a field that does not allow carriage returns (for example, the typical Name field) and press Enter. In many cases, this is the equivalent of clicking OK.

  – **Cancel**—Click the X on the right side of the window's title bar.

  – **Help**—Press F1.

Accessing Online Help

To access online help for Security Manager, do one of the following:

- To open the main Security Manager online help page, select **Help > Help Topics**.
- To open context-sensitive online help for the active page, select **Help > Help About This Page** or click the **?** button in the toolbar.
- To open context-sensitive online help for a dialog box, click **Help** in the dialog box.

Tip You must configure Internet Explorer to allow active content to run on your computer for the online help to open unblocked. In Internet Explorer, select **Tools > Internet Options** and click the **Advanced** tab. Scroll to the Security section, and select **Allow active content to run in files on My Computer**. Click **OK** to save the change. For a complete list of configuration requirements for Internet Explorer and Firefox browsers, see the *Installation Guide for Cisco Security Manager*.

The online help page appears without any user authentication. Though the pages are opening with direct URL access, they are only static content pages and function within the Cisco Security Manager.

Preparing Devices for Management

Before you start to manage a device using Security Manager, you should prepare the device with at least a minimal configuration. The following sections describe the basic device configurations needed for various transport protocols or device types. Before configuring transport protocols, determine the requirements for your devices by reading Understanding Device Communication Requirements, page 2-1.

- Understanding Device Communication Requirements, page 2-1
- Setting Up SSL (HTTPS), page 2-3
- Setting Up SSH, page 2-5
- Setting Up AUS or Configuration Engine, page 2-8
- Configuring Licenses on Cisco ASA Devices, page 2-9
- Configuring Licenses on Cisco IOS Devices, page 2-10
- Initializing IPS Devices, page 2-10

Understanding Device Communication Requirements

Security Manager provides many different ways for you to manage devices. The easiest methods involve Security Manager directly contacting the devices. Security Manager might access a device during inventory or policy discovery, during configuration deployment, or in response to actions you take in Security Manager that request device contact (such as testing connectivity).

Because you can use off-line methods to add devices to the Security Manager inventory or to deploy configuration changes to the devices, configuring device communication settings for Security Manager's use is optional. However, you typically need to configure basic device communication settings on the devices to implement your off-line or customized configuration deployment tools.

In Security Manager, you can configure which transport protocol to use as the default for a type of device, and change it for specific devices that are configured to respond to a different protocol. Security Manager is configured with default protocols that are the most commonly-used protocols for that type of device. To change the default device communication setting for a type of device, select **Tools > Security Manager Administration** and select **Device Communication** from the table of contents (for more information, see Device Communication Page, page 11-21). To change the transport setting for a specific device, modify its device properties as described in Viewing or Changing Device Properties, page 3-40.

Security Manager can use these transport protocols:

- SSL (HTTPS)—Secure Socket Layer, which is an HTTPS connection, is the only transport protocol used with PIX Firewalls, Adaptive Security Appliances (ASA), and Firewall Services Modules (FWSM). It is also the default protocol for IPS devices and for routers running Cisco IOS Software release 12.3 or higher.

  If you use SSL as the transport protocol on Cisco IOS routers, you must also configure SSH on the routers. Security Manager uses SSH connections to handle interactive command deployments during SSL deployments.

  Cisco Security Manager was using OpenSSL for the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols. Beginning with version 4.13, Cisco Security Manager replaced OpenSSL version 1.0.2 with Cisco SSL version 6.x. Cisco SSL enables FIPS compliance over full FIPS Validation which results in fast and cost-effective connectivity. The Common Criteria mode in Cisco SSL allows easier compliance. Cisco SSL is feature-forward when compared to OpenSSL. The product Security Baseline (PSB) requirements for CiscoSSL ensures important security aspects such as credential and key management, cryptography standards, antispoofing capabilities, integrity and tamper protection, and session, data, and stream management and administration are taken care of.

  For information on configuring SSL, see Setting Up SSL (HTTPS), page 2-3.

- SSH—Secure Shell is the default transport protocol for Catalyst switches and Catalyst 6500/7600 devices. You can also use it with Cisco IOS routers.

  For information on configuring SSH, see Setting Up SSH, page 2-5.

- Telnet—Telnet is the default protocol for routers running Cisco IOS software releases 12.1 and 12.2. You can also use it with Catalyst switches, Catalyst 6500/7600 devices, and routers running Cisco IOS Software release 12.3 and higher. See the Cisco IOS software documentation for configuring Telnet.

- HTTP—You can use HTTP instead of HTTPS (SSL) with IPS devices. HTTP is not the default protocol for any device type.

- TMS—Token Management Server is treated like a transport protocol in Security Manager, but it is not a real transport protocol. Instead, by configuring TMS as the transport protocol of a router, you are telling Security Manager to deploy configurations to a TMS. From the TMS, you can download the configuration to an eToken, plug the eToken into the router's USB bus, and update the configuration. TMS is available only for certain routers running Cisco IOS Software 12.3 or higher.

  For information on deploying configurations to a TMS and downloading them to a router, see Deploying Configurations to a Token Management Server, page 8-43.

Security Manager can also use indirect methods to deploy configurations to devices, staging the configuration on a server that manages the deployment to the devices. These indirect methods also allow you to use dynamic IP addresses on your devices. The methods are not treated as transport protocols, but as adjuncts to the transport protocol for the device. You can use these indirect methods:

- AUS (Auto Update Server)—When you add a device to Security Manager, you can select the AUS server that is managing it. You can use AUS with PIX Firewalls and ASA devices.

  For information on configuring a device to use an AUS server, see Setting Up AUS or Configuration Engine, page 2-8.

- Configuration Engine—When you add a router to Security Manager, you can select the Configuration Engine that is managing it.

  For more information on configuring a router to use a Configuration Engine server, see Setting Up AUS or Configuration Engine, page 2-8.

For information on adding devices that use AUS or Configuration Engine servers to Security Manager, and how to add the servers, see these topics:

- Adding Devices to the Device Inventory, page 3-6
- Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36

Setting Up SSL (HTTPS)

With many devices, you can use the Secure Socket Layer (SSL) protocol, also known as HTTPS, to communicate with the device. When you deploy configurations with this protocol, Security Manager encrypts the configuration file before sending it to the device.

The following topics describe how to set up SSL on the devices:

- Setting Up SSL (HTTPS) on PIX Firewall, ASA and FWSM Devices, page 2-3
- Setting Up SSL on Cisco IOS Routers, page 2-4

Setting Up SSL (HTTPS) on PIX Firewall, ASA and FWSM Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

This procedure describes the tasks to complete before you use SSL as the transport protocol for device management on PIX Firewall, ASA and FWSM devices.

Step 1 Enter configuration mode.

```
hostname# config terminal
```

Respond to the prompts appropriately. Here are some tips:

- Enter **y** when the prompt asks if you want to preconfigure using interactive prompts.
- Enter the current enable password.
- Specify the time zone, year, month, day, and time.
- If the device:
  - Is new—Specify the network interface IP address and network mask that applies to the inside IP address of the device.
  - Exists—Verify that the interface IP address and mask are correct.
- If the device:
  - Is new—Specify the hostname and the domain name.
  - Exists—Verify that the hostname and domain name are correct.
- When prompted for the IP address of the host that runs the PIX Device Manager, specify the IP address of the Security Manager server.
- Enter **yes** when the prompt asks if you want to write the above changes to Flash.

Step 2 If you are configuring an ASA, specify the SSL/TLS protocol version the ASA uses when acting as a server. Beginning with version 4.8, Security Manager supports all SSL/TLS protocol versions, the latest certified version being TLS 1.2.

```
hostname(config)# ssl server-version any
```

Step 3 Enable the HTTP server.

```
hostname(config)# http server enable
```

Step 4 Specify the host or network authorized to initiate an HTTP connection to the device.

```
hostname(config)# http ip_address  [netmask] [if_name]
```

Where:

- *ip_address*—The IP address of the Security Manager server.
- *netmask*—The network mask for the IP address.
- *if_name*—The device interface name (default is **inside**) from which Security Manager initiates the HTTP connection.

Step 5 Save the current configuration in Flash memory.

```
hostname(config)# write memory
```

Setting Up SSL on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS it does not support any enhancements.

This procedure describes the tasks to complete before you use SSL as the transport protocol for device management on Cisco IOS routers.

Step 1 Enter configuration mode.

```
hostname# config terminal
```

Step 2 Configure the hostname and domain name if the device is new.

```
router(config)# hostname name
hostname(config)# ip domain-name your_domain
```

Step 3 Configure level 15 privilege. SSL requires that you must have level 15 privileges to log in to a Cisco IOS router.

```
hostname(config)# username username privilege 15 password 0 password
```

Step 4 Enable either local authorization or AAA authorization:

- Local authorization— If you are using AAA for authorization but would like to use local authorization, use the following commands to disable AAA authorization and AAA authentication at login, where *list-name* is a character string used to name the list of authorization methods, and to enable local authorization using the username you just configured:

```
hostname(config)# no aaa authorization network list-name
hostname(config)# no aaa authentication login list-name
hostname(config)#ip http authentication local
```

If you do not enter the **ip http authentication local** command, the default enable password is used for authentication.

- AAA authorization—Use the following commands to enable AAA authentication and authorization. The last two commands are necessary only if multiple AAA lists are defined; *list-name* is a character string used to name the list of authorization methods. These commands authenticate the user that is contacting the device using the HTTPS protocol.

```
hostname(config)#ip http authentication aaa
hostname(config)#ip http authentication aaa login-authentication list-name
hostname(config)# ip http authentication aaa exec-authorization list-name
```

Step 5 Enable the HTTPS server.

```
hostname(config)# ip http secure-server
```

Step 6 Exit configuration mode and return to Exec mode.

```
hostname(config)# exit
```

Step 7 Verify that SSL is set up on the device. The Device should respond with an "enabled" status.

```
hostname# show ip http server secure status
```

Setting Up SSH

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any enhancements.

You can use the Secure Shell (SSH) protocol to communicate with Cisco IOS Routers, Catalyst switches, and Catalyst 6500/7600 devices. This protocol provides strong authentication and secure communications over insecure channels. Security Manager supports both SSH versions 1.5 and 2. Once connected to the device, Security Manager determines which version to use and communicates using that version.

The following topics describe how to set up SSH on the supported devices:

- Critical Line-Ending Conventions for SSH, page 2-5
- Testing Authentication, page 2-6
- Setting Up SSH on Cisco IOS Routers, Catalyst Switches, and Catalyst 6500/7600 devices, page 2-6
- Preventing Non-SSH Connections (Optional), page 2-7

Critical Line-Ending Conventions for SSH

The following line-ending conventions for SSH must be observed to avoid system failure:

- Do not end banner message lines with "#", "# ", ">", or "> ". If your system requires a pound sign or greater-than sign at the end of a banner message, ensure that it is followed by two spaces.
- Do not use banner message lines that contain only "Username: " or "Password: "
- Do not customize the device user EXEC mode prompt to not end with ">" or "#".

Testing Authentication

Before you set up SSH, you must test authentication without SSH to make sure the device can be authenticated. You can authenticate with a local username and password or with an authentication, authorization, and accounting (AAA) server running TACACS+ or RADIUS.

This procedure describes how to test authentication without SSH using a local or AAA server username and password.

Step 1 Enter configuration mode.

```
router# config terminal
```

Step 2 Specify that the local username and password should be used in the absence of AAA statements. On Cisco IOS routers, you can use the **login local** command on VTY lines instead of the **aaa new-model** command.

```
hostname(config)#aaa new-model
```

Step 3 (Optional) Configure a user account in the local database of the device.

```
hostname(config)# username name password 0 password
```

Step 4 Exit configuration mode and return to Exec mode.

```
hostname(config)# exit
```

Step 5 Save the configuration changes.

```
hostname(config)# write memory
```

Setting Up SSH on Cisco IOS Routers, Catalyst Switches, and Catalyst 6500/7600 devices

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS it does not support any enhancements.

This procedure describes the tasks required to set up SSH on Cisco IOS routers, Catalyst switches, and Catalyst 6500/7600 devices.

Tip You must configure SSH on Cisco IOS routers because Security Manager uses SSH connections to handle interactive command deployments during SSL deployments.

Related Topics

- Critical Line-Ending Conventions for SSH, page 2-5

- Testing Authentication, page 2-6
- Preventing Non-SSH Connections (Optional), page 2-7

Step 1 Enter configuration mode.

```
router# config terminal
```

Step 2 Configure the hostname and domain name if the device is new.

```
router(config)# hostname name
hostname(config)# ip domain-name your_domain
```

Step 3 Generate the RSA key pair for the SSH session. When the device prompts you to enter the size of the modulus, we recommend that you enter 1024.

```
hostname(config)# crypto key generate rsa
```

Step 4 (Optional) Set the timeout interval in minutes and the number of retries.

```
hostname(config)# ip ssh timeout time
hostname(config)# ip ssh authentication-retries n
```

Step 5 Exit configuration mode and return to Exec mode.

```
hostname(config)# exit
```

Step 6 Save the configuration changes.

```
hostname# write memory
```

Preventing Non-SSH Connections (Optional)

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS it does not support any enhancements.

After configuring SSH, you can configure the Cisco IOS routers, Catalyst switches, and Catalyst 6500/7600 devices to use SSH connections only.

Related Topics

- Critical Line-Ending Conventions for SSH, page 2-5
- Testing Authentication, page 2-6
- Setting Up SSH on Cisco IOS Routers, Catalyst Switches, and Catalyst 6500/7600 devices, page 2-6

Step 1 Enter configuration mode.

```
router# config terminal
```

Step 2 Set up the router for Telnet access, specifying the first and last line numbers that can be used (numbers range from 0 to 1180, and the last number must be greater than the first number).

```
hostname(config)# line vty first_line last_line
```

Step 3 Prevent non-SSH connections, such as Telnet.

```
hostname(config-line)# transport input ssh
```

Step 4 Exit configuration mode.

```
hostname(config-line)# end
```

Step 5 Save the configuration changes.

```
hostname# write memory
```

Setting Up AUS or Configuration Engine

With many devices, you can use an intermediate transport server to stage configuration updates to the device. These transport servers can also allow you to manage devices that use dynamically assigned IP address (using a DHCP server) instead of static IP addresses. When you deploy configurations using a transport server, Security Manager deploys the configuration to the server, and the device retrieves the configuration from the server. You can use Auto Update Server, running the AUS protocol, or Cisco Configuration Engine, running the CNS protocol.

The following topics describe how to set up AUS or CNS on the devices:

- Setting Up AUS on PIX Firewall and ASA Devices, page 2-8

Setting Up AUS on PIX Firewall and ASA Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

You can configure PIX firewalls and ASA devices to use the AUS protocol to contact an Auto Update Server or CNS Configuration Engine for configuration and image updates. When using Configuration Engine, the device uses the same AUS protocol used for Auto Update Server, so the configuration is the same. For an end-to-end explanation of how AUS/CE deployment works, see Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41.

You need to initially configure AUS settings on the device so that the device knows that it must contact the AUS/CE server for configuration updates. After the initial deployment, you can change these settings using the **Platform > Device Admin > Server Access > AUS** policy.

This procedure describes the tasks to complete before you use AUS or CNS as the transport protocol for device management on PIX firewall and ASA devices.

Step 1 Enter configuration mode.

```
router# config terminal
```

Step 2 Connect to the AUS. Specify a username and its password that can log into Security Manager. The port number is typically 443.

```
hostname(config)# auto-update server https:// username:password@AUSserver_IP_address:port
/autoupdate/AutoUpdateServlet
```

Step 3 Specify the polling period for AUS.

```
hostname(config)# auto-update poll-period  poll_period  [retry_count ] [retry_period ]
```

Where:

- *poll_period*—The polling period interval between two updates. Default is 720 minutes (12 hours).
- *retry_count*—(Optional) The number of times to retry if the server connection attempt fails. Default is 0.
- *retry_period*—(Optional) The number of minutes between retries. Default is 5.

Step 4 Configure the device to use the specified unique device ID to identify itself.

```
hostname(config)# auto-update device-id  [ hardware-serial | hostname |
ipaddress  [if_name ] |  mac-address [if_name ] | string text ]
```

Where:

- *if_name*—The device interface name (the default is **inside**).
- *text*—A unique string name.

Step 5 Save the configuration changes.

```
hostname# write memory
```

Configuring Licenses on Cisco ASA Devices

Devices that run Cisco ASA Software require Product Activation Keys for each feature license. Some licenses are optional, such as Botnet Traffic Filtering, and can be time-based. Other features are standard on some models, but optional on others, such as the Failover license, which is optional on the 5505 and 5510 models but standard on all other models.

You cannot install or activate ASA licenses through Security Manager. Instead, use the Adaptive Security Device Manager (ASDM). Enter the activation keys by selecting **Configuration > Device Management > Licensing > Activation Key** and following the instructions in the online help for that page. The Activation Key page also lists the state of all feature licenses. The ASDM online help includes extensive information about ASA licensing.

When you deploy configurations from Security Manager, the device must have active licenses for all features in the configuration or you will see deployment errors. In most cases, Security Manager does not prevent you from configuring a feature based on the licenses that are active on a device. For example, you can configure Botnet Traffic Filtering for a device even if that device has a disabled Botnet license.

The exception is the Failover license on the 5505 and 5510 models. There is a device property that you can set to indicate whether there is an active Failover license on a device: License Supports Failover. You can set this property by double-clicking the device (in Device view) to open the Device Properties page; the option is on the General tab (see Device Properties: General Page, page 3-41). If you discover policies on the device, for example, when adding the device to the inventory using the Add Device From Network or Add Device from File (from an inventory file, not a configuration file) options, Security Manager determines the state of the Failover license and sets the property appropriately. You are responsible for ensuring that the property remains accurate. You will see deployment failures if the property is selected but the device has an inactive Failover license.

Tip If you add the device using the New Device or Configuration File options, you can set the License Supports Failover property while adding the device instead of waiting to set it in the device properties.

Configuring Licenses on Cisco IOS Devices

Devices that run Cisco IOS Software require license files for various features, including security features. If these licenses are not installed on the device (such as the securityk9 package), Security Manager cannot configure commands that require a particular license level, and you will experience deployment failures when you try to deploy your policies to an unlicensed device.

Although you can use Security Manager to deploy and manage IPS licenses, you cannot use it to deploy and manage any other type of license. Configure these licenses directly on the device using the command line interface or use Cisco License Manager. Following is the general process for configuring licenses. For more information about configuring licenses, see Cisco IOS Software Activation Command Guide and Cisco IOS Software Activation Command Reference on Cisco.com.

1. Obtain the licenses required for the features you want to use or you can use the evaluation licenses that come bundled with some devices. Use the **show license all** command to view the available licenses.
2. Copy the purchased licenses to the flash storage on the device or put them on a TFTP server. For example, you could place the licenses on a TFTP server and use the **copy tftp flash0:** command to copy the files to the flash0 storage area.
3. Use the **license install** command to install each purchased license. For example:

   license install flash0:uc-base-CISCO2951-FHH1216P06Z.xml

   Some licenses prompt you to read and accept a license agreement.

   If you want to use an evaluation license, use the **license boot** command to enable them and then reload the device. You must accept the end-user license agreement before Security Manager can deploy configurations to the device.
4. You can use the **show version**, **show license feature**, and **show license all** commands to check on your installed licenses.

Initializing IPS Devices

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

To initialize an IPS device, you must configure the following settings. These are network settings, and only a user with administrator privileges on the IPS device can configure them:

- Sensor name
- IP address
- Netmask
- Default route
- Enable TLS/SSL (to enable TLS/SSL in the web server on the device)
- Web server port
- Use default ports

You configure these settings through the **setup** command in Intrusion Prevention System Device Manager (IDM) or in a command-line session, depending upon which platform is used by your IPS device. For a list of supported IPS platforms, see the supported devices and software versions information at the following URL: http://www.cisco.com/en/US/products/ps6498/products_device_support_tables_list.html.

For detailed information on these settings, refer to the technical documentation for your IPS device.

Note For information on preparing an IOS IPS device for use, see Initial Preparation of a Cisco IOS IPS Router, page 45-5.

Managing the Device Inventory

Before you can manage devices in Security Manager, you must prepare the devices for management, then add those devices to the Security Manager device inventory. After you add the devices, you can view and edit device information, configure policies on devices, copy and share policies, clone devices, and so on. The following topics describe how to manage the device inventory:

- Understanding the Device Inventory, page 3-1
- Adding Devices to the Device Inventory, page 3-7
- Working with the Device Inventory, page 3-36
- Working with Device Groups, page 3-59
- Working with Device Status View, page 3-64

Understanding the Device Inventory

Security Manager maintains an inventory of the devices that it manages. The inventory includes the information required to locate and log into the device, so that your policies can be deployed to the devices. The following topics describe some general concepts related to the device inventory:

- Understanding the Device View, page 3-1
- Understanding Device Names and What Is Considered a Device, page 3-3
- Understanding Device Credentials, page 3-4
- Understanding Device Properties, page 3-6

Understanding the Device View

The Device View button opens the Devices page, from which you can add and delete devices from the Security Manager inventory and manage device policies, properties, and interfaces centrally.

This is a device-centric view in which you can see all devices that you are managing and you can select specific devices to view their properties and define their settings and policies. You can define security policies locally on specific devices. You can then share those policies to make them globally available to be assigned to other devices.

The Devices page contains two panes. The left pane contains two elements: the Device selector, located in the top left pane, and the Policy selector, located in the bottom left pane. The right pane is the main content area. The following illustration shows the Devices page.

Figure 3-1 Devices Page

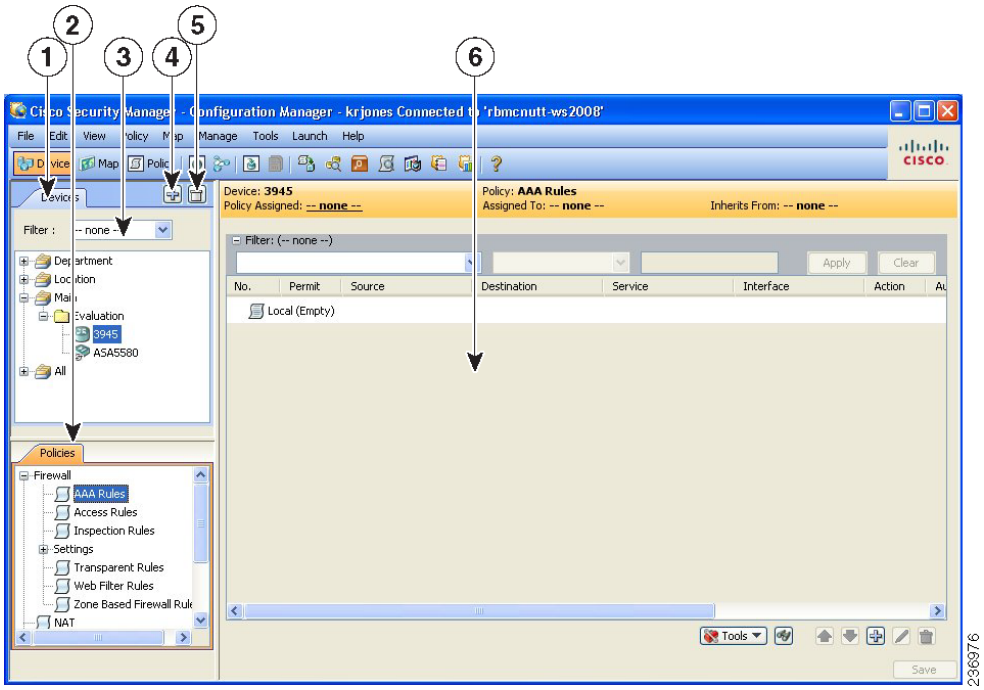

Device selector (1, 3, 4, 5)—Contains the following:

- Add and Delete buttons (4, 5)—Enables you to add and delete devices from the Security Manager inventory.
- Filter field (3)—Enables you to display a subset of devices based on the filtering criteria you define. For details, see Filtering Items in Selectors, page 1-45.
- Device tree—Lists the device groups and devices that exist in the system. Each device type is represented by an icon. For information about the icons, see Figure 3-2.

  If you hover the mouse pointer over a device, detailed information about the device appears in a popup window. The information is a summary of the device properties (see Device Properties: General Page, page 3-41).

Note Beginning from version 4.8, Security Manager displays the updated version information of a device that has been upgraded using Auto Update Server (AUS). To enable this feature you must configure Security Manager details in the AUS user interface. If you hover the mouse over a device, the following message appears if AUS has successfully updated the device version:

"State Description: Version update is successfully completed by Auto Update Server. Check if any other configuration changes are required in Security Manager."

Figure 3-2 Device Icons

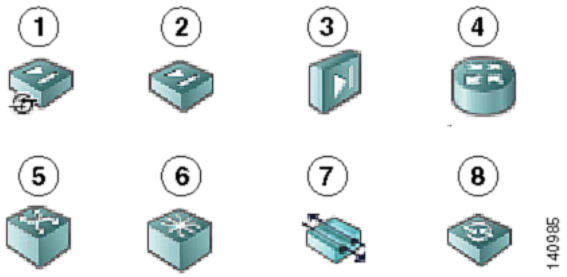

1	Adaptive Security Appliances (ASA)	**5**	Catalyst Switch
2	PIX Firewall	**6**	Catalyst 7600 Series Router
3	Catalyst security Services Modules: Firewall Services Module (FWSM) and ASA-SM	**7**	VPN 3000 Concentrator
4	Cisco IOS Router	**8**	Intrusion Prevention System (IPS)

- Shortcut menu options—When you right-click a device or device group, you get a menu of commands related to that device or group. These commands are shortcuts to commands available in the regular menus.

Policy selector (2)—Contains the following:

- Policy groups—Lists the policy groups that are supported on the selected device type. The policy groups that are displayed are dependent on four factors:
  - The type of device selected in the Device selector.
  - The operating system running on the device.
  - The target operating system version selected for determining which commands will be available for generated configurations.
  - Whether the device contains supported service modules.

  For more information about policies, see Understanding Policies, page 5-1.
- Shortcut menu options—When you right-click a policy, you get a menu of commands related to that policy. These commands are shortcuts to commands available in the regular menus.

Contents pane (6)—The main content area.

The information displayed in this area depends on the device you select from the Device selector and the option you select from the Policy selector.

Understanding Device Names and What Is Considered a Device

Besides managing traditional devices, you can use Security Manager to manage virtual devices that you can define on some types of security devices. These virtual devices are treated as separate devices in the device inventory, and they appear as separate entries in the device selectors. Because these virtual devices actually reside on a host physical device, many actions, such as deployment, will have to include the host device as well as the virtual device.

All physical devices appear in the device selectors. In addition, these are the types of virtual devices that appear in the device selectors:

- Security Contexts—You can define security contexts on PIX Firewall, FWSM, and ASA devices. Security contexts act as virtual firewalls. By default, security contexts appear in the device selectors using this naming convention: *host-display-name_context-name*, where *host-display-name* is the display name of the device on which the context is defined, and *context-name* is the name of the security context. For example, the admin security context on the device named firewall12 would be called firewall12_admin.

Tip You can control whether the display name is added to the context name using the **Prepend Device Name when Generating Security Context Names** property on the Discovery settings page (see Discovery Page, page 11-25). However, if you do not add the display name, it is very difficult to determine the hosting device for a context, and the context names are not sorted with the host device (they do not appear in a folder attached to the host device). If you do not add the display name, Security Manager adds a numeric suffix to the context name if more than one context of the same name is added to the inventory (for example, admin_01, admin_02), and these numbers are not related to the host device.

- Virtual Sensors—You can define virtual sensors on IPS devices. Virtual sensors appear in device selectors using the *host-display-name_virtual-sensor-name* naming convention, and there is not a discovery setting to control this convention.

Tip You can always change the display name for a virtual sensor, security context, or other type of device in the device's properties.

Besides the naming conventions for virtual devices, you also need to understand the relationship between various types of device names:

- Display name—The display name is simply the name that appears within Security Manager in device selectors. This name does not have to be related to any name actually defined on the device. When you add devices to the inventory, a display name is suggested based on the DNS name or IP address you enter, but you can use whatever naming convention you want to use.
- DNS name—The DNS name you define for a device must be resolvable by the DNS server configured for the Security Manager server.
- IP address—The IP address you define for a device should be the management IP address for the device.
- Hostname—When you discover a device, the hostname property that is shown in the device properties is taken from the device's configuration. If you add devices using configuration files, and a file does not contain a hostname command, the initial hostname is the name of the configuration file.

  However, the hostname device property is not updated if you change the hostname on the device. There is a Hostname policy in the device platform policy area, and it is this Hostname policy that determines the hostname that is defined on the device.

Understanding Device Credentials

Security Manager requires credentials for logging in to devices. You can provide device credentials in two ways:

- When you add a device manually or from network discovery. For more information, see these topics:
  - Adding Devices from the Network, page 3-12
  - Adding Devices by Manual Definition, page 3-26
- By editing the device properties. For more information, see Viewing or Changing Device Properties, page 3-41.

You can provide the following device credentials:

- Primary Credentials—The username and password for logging into the device using SSH or Telnet. This information is required for device communication.
- HTTP Credentials—Some devices allow HTTP or HTTPS connections, and some devices (such as IPS devices) require it. By default, Security Manager uses the primary credentials for HTTP/HTTPS access, but you can configure unique HTTP/HTTPS credentials.
- RX-Boot Mode—(Optional) Some Cisco routers are designed to run from flash memory where they boot only from the first file in flash. This means that you must run an image other than the one in flash to upgrade the flash image. That image is a reduced command-set image referred to as RX-Boot (a ROM-based image).
- SNMP Credentials—(Optional) The Simple Network Management Protocol (SNMP) facilitates the exchange of management information between network devices. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

Note PIX/ASA/FWSM devices require that user names be at least four characters. Passwords can be three to 32 characters; we recommend that passwords be at least eight characters. For ASA devices running the software version 9.6(1) or later, you can enter a password up to 127 characters.

Rather than using device-based credentials, you can configure Security Manager to use the credentials you use when you log into Security Manager. You can then use the AAA server's accounting facilities to track configuration changes by user. Using user login credentials is suitable only if your environment is configured according to these standards:

- You use TACACS+ or RADIUS for change auditing. User-login credentials will be reflected in these accounting records. If you use device credentials, all changes made through Security Manager will come from the same account, regardless of which user made the change.
- User accounts are configured in the AAA server, and they have appropriate device-level access to perform configuration changes.
- You configure Security Manager and the managed devices to use the AAA server for authorization. For information on configuring Security Manager to use AAA, see the *Installation Guide for Cisco Security Manager*.
- You do not use one-time passwords.

If your network setup supports using user-login credentials, you can configure Security Manager to use them by selecting **Tools > Security Manager Administration**. Select **Device Communication** from the table of contents, and select **Security Manager User Login Credentials** in the **Connect to Device Using** field. The default is to use device credentials for all device access.

Related Topics

- Device Credentials Page, page 3-45
- Adding Devices to the Device Inventory, page 3-7

- Device Communication Page, page 11-21

Understanding Device Properties

You define device properties when you add devices to Security Manager. Device properties are general information about the device, credentials, the group the device is assigned to, and policy overrides. You must provide some device property information, such as device identity and primary credentials, when you add the device, but you can add or edit the properties from the Device Properties dialog box.

To view the device properties, do one of the following in the Device selector:

- Double-click a device.
- Right-click a device and select **Device Properties**.
- Select a device and select **Tools > Device Properties**.

The Device Properties dialog box has two panes. The left pane contains a table of contents with these items:

- General—Contains general information about the device, such as device identity, the operating system running on the device, and device communication settings.
- Credentials—Contains device primary credentials (username, password, and enable password), SNMP credentials, Rx-Boot Mode credentials, and HTTP credentials.
- Device Groups—Contains the groups to which the device is assigned.
- Cluster Information—Contains detailed informatio for the cluster group, if any.
- Policy Object Overrides—Contains global settings of certain types of reusable policy objects that you can override for this device.

When you select an item in the table of contents, the corresponding information is displayed in the right pane.

Notes

- Security Manager does not assume that the DNS hostname that appears on the Device Properties page is the same as the hostname that you configured on the device.
- When you add a device to Security Manager, you must enter either the management IP address or the DNS hostname. Because it is not possible to determine the management interface and, therefore, the management IP address when you discover from a configuration file, the hostname in the configuration file is used as the DNS hostname. If the hostname is missing in the CLI of the configuration file, the configuration filename is used as the DNS hostname.
- When you discover a device from the network, the DNS hostname in the Device Properties page is not updated with the hostname configured on the device. Therefore, if you want to specify the DNS hostname for the device, you must specify it manually when you add the device to Security Manager or on the Device Properties page.

For more information about device properties, see Viewing or Changing Device Properties, page 3-41.

Adding Devices to the Device Inventory

When you add a device to Security Manager, you specify the identifying information for the device, such as its DNS name and IP address. This information is added during device discovery. You can also bring in existing network configurations associated with a device by initiating policy discovery. For complete information on policy discovery, see Discovering Policies, page 5-12. Once you add the device, it appears in the Security Manager device inventory.

The New Device wizard guides you through the process of adding devices to the inventory. You can add devices from many different sources, and the path through the wizard differs significantly based on the method you are using.

To start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.

Note There is also another way to add devices. If you exported a .dev file from another Security Manager server, which contains not only a device inventory but also the policies and policy objects assigned to them, you can import the file using the **File > Import** command. For more information, see Importing Policies or Devices, page 10-13.

Tips on Adding Devices and Service Modules

- For PIX Firewalls and FWSM and ASA devices that are configured for failover, add only the active unit to Security Manager. Ensure that the device is configured with a management IP address and use that address for discovery. When discovering Catalyst switches that contain more than one service module (FWSM or ASA-SM) configured for failover, when prompted, select **Do Not Discover Module** for the failover modules. Security Manager always manages the active admin context, regardless of whether you added the primary or secondary failover service module.
- Security Manager can manage ASA clusters after they have been configured as a cluster using the CLI bootstrapping as defined in the ASA Configuration Guide (see http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html). All the members of a cluster are assigned individual IP addresses during the bootstrap process. When adding a cluster to Security Manager, you do so by discovering the cluster using the main cluster IP address. The main cluster IP address is a fixed address for the cluster that always belongs to the current control unit. This is not the control unit's individual IP Address. For more information about clusters, see Working with Device Clusters, page 3-9.
- Service modules are treated as separate devices. For most modules, you must add the service module separately from its host device. However, Security Manager can automatically discover FWSM or IDSM modules in a Catalyst 6500 device, so you need only add the parent device. (You cannot discover an ASA-SM during discovery of the parent device. You must add the ASA-SM separately.) The only exception is if you configure an FWSM or IDSM module to use a non-default port for HTTPS (SSL), in which case you must add the module separately.
- When adding an ASA-SM or FWSM that has multiple security contexts (they are running in multiple-context mode), do not add the security contexts individually using their management IP addresses. Instead, add the device using the admin context management address (this also adds the individual contexts). Then, configure Security Manager to deploy configurations to multiple-context devices serially as described in Changing How Security Manager Deploys Configurations to Multiple-Context FWSM, page 9-17.
- You cannot add devices beyond the device limits defined by your Security Manager license. For example, if you have a license for 50 devices, and there are 45 devices in the inventory, if you try to add a multiple-context ASA with 6 security contents, the device addition and discovery fails.

The following topics describe the various methods of adding devices:

- **Add Device from Network**—To add devices that are currently active on the network, see Adding Devices from the Network, page 3-12. Security Manager connects directly and securely to the device and discovers its identifying information and properties.
  - **Pros**—You need to specify minimal information about a device, and Security Manager obtains the detailed information directly from the device, ensuring accuracy.
  - **Cons**—You can add only one device at a time. You cannot add devices that have dynamic IP addresses, unless you determine the device's current IP address, add it using that address, and then update the device properties in Security Manager to identify the Configuration Engine that is managing the device.
- **Add from Configuration File**—To add devices by using a copy of the device configuration files, see Adding Devices from Configuration Files, page 3-22.
  - **Pros**—You can add more than one device at a time.
  - **Cons**—You cannot use this method to add Catalyst 6500/7600 or IPS devices. When adding groups of configuration files, all files must be for the same device type.

    Also, you cannot successfully discover policies that require a connection with the device. For example, if a policy points to a file that resides on the device, adding the device using the configuration file will result in a Security Manager configuration that includes the **no** form of the command, because Security Manager cannot retrieve the referenced file from the device. For example, the **svc image** command for web VPNs might be negated.
- **Add New Device**—To add a device that does not yet exist in the network, so that you can pre-provision it in Security Manager, see Adding Devices by Manual Definition, page 3-26. You can create the device in the system, assign policies to the device, and generate configuration files before installing the device hardware.
  - **Pros**—You can pre-provision devices that do not yet exist in the network.
  - **Cons**—You must specify more information than that required by any other method. If you create a Catalyst 6500 device, or a router that contains an IPS module, you should discover its modules by selecting **Policy > Discover Policies on Device**.
- **Add Device from File**—To add devices from an inventory file in comma-separated values (CSV) format, see Adding Devices from an Inventory File, page 3-31.
  - **Pros**—You can add multiple devices of different types at one time. You can reuse the inventory list from your other network management applications, including CiscoWorks Common Services, Cisco Security Monitoring, Analysis and Response System (CS-MARS), and other Security Manager servers. If you use a file exported from another Security Manager server, you can optionally add the devices without discovering policies, which is convenient for adding offline or standby devices.
  - **Cons**—You cannot use this method to update the properties of devices already defined in the inventory. Also, policy discovery can fail if you attempt to import more than 100 devices at one time, and might fail for even fewer devices. In the case of IPS devices, do not add more than four IPS devices at a time to avoid policy discovery failures.

Working with Generically Supported Devices

Security Manager can manage some device models even if the model does not appear in the supported device list. This type of generic device support relies on the fact that device features are controlled more by the software running on the device than the device model. If you have a device that does not appear in the explicitly supported device list, you can try to manage it as a generic device using the device modules listed in the table below.

You can add a generic router using any of the various methods for adding devices:

- If you are adding a router directly from the network that is not natively supported and can be generically supported, Security Manager automatically discovers the device as a generic router. You will be prompted whether you want to add the generic device or cancel the operation.
- If you are adding a generic router by manual definition or from a configuration file, you must select the appropriate generic device type according to the OS running on the router (see Table 3-1 on page 3-9).
- If you are adding a generic router from an inventory file, you must make sure that you use the correct SYSOBJID in the CSV file. Use 1.0.0.0.0.0.0.1 for devices running Cisco IOS Software (Cisco Generic Integrated Services Router) and use 1.0.0.0.0.0.0.2 for devices running Cisco IOS XE Software (Cisco Generic Aggregation Services Router).

Tips

- Generic router support encompasses ISR and ASR hardware models that are not natively supported. It does not encompass software versions that are not yet supported. In other words, Security Manager allows you to manage unsupported hardware platforms if those platforms are running a supported software release. This type of generic support works best for new models of series that are already explicitly supported. For example, a new model in the ASR 1000 series, or in the ISR 88x or 89x series. Generic support does not work with carrier-class routers (the CRS) or for Catalyst switches.
- Because this support is generic, Security Manager cannot determine if a particular feature is not available on the specific model you are managing. You are responsible for determining if a feature that you are allowed to configure in Security Manager is supported on the device. If you configure an unsupported feature, you will see errors when you deploy the configuration to the device.
- If the device contains an explicitly supported module, such as an AIM-IPS module, the module is also supported. However, the module's model must be explicitly supported: there is no generic module support.

Table 3-1 Cisco Security Manager Generically Supported Devices

Generic Device Type	When To Use
Cisco Generic Aggregation Services Router	For devices running Cisco IOS XE Software.
Cisco Generic Integration Services Router	For devices running Cisco IOS Software.

Working with Device Clusters

Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. Clustering is supported on ASA 5580 and 5585 devices running 9.0(1) or later and on ASA 5512-X, 5515-X, 5525-X, 5545-X and 5555-X devices running 9.1(4) or later.

Security Manager can manage ASA clusters after they have been configured as a cluster using the CLI bootstrapping as defined in the ASA Configuration Guide (see http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html).

All the members of a cluster are assigned individual IP addresses during the bootstrap process. When adding a cluster to Security Manager, you do so by discovering the cluster using the main cluster IP address. The main cluster IP address is a fixed address for the cluster that always belongs to the current control unit. This is not the control unit's individual IP Address.

Note You cannot convert a standalone device to a cluster in Security Manager by rediscovering the device after performing the necessary CLI bootstrapping. You must first delete the device from Security Manager, and then after performing the necessary CLI bootstrapping, you can add the cluster to Security Manager as a new device.

The cluster is represented as a single device in Security Manager. After the cluster has been added to Security Manager, you can finish configuring the cluster settings such as cluster interfaces and security policies.

Note Clustering has specific configuration requirements and restrictions. Please refer to the ASA documentation at http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html for detailed information about requirements, configuration recommendations, and performance information.

Unsupported Features on ASA Clusters

These features cannot be configured with clustering enabled, and the commands will be rejected.

- Unified Communications
- Remote access VPN (SSL VPN and IPsec VPN)
- The following application inspections:
  - CTIQBE
  - GTP
  - H323, H225, and RAS
  - IPsec passthrough
  - MGCP
  - MMP
  - RTSP
  - SIP
  - SCCP (Skinny)
  - WAAS
  - WCCP
- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, relay, and proxy

- VPN load balancing
- Failover
- ASA CX module

Centralized Features

The following features are only supported on the control unit, and are not scaled for the cluster. For example, you have a cluster of eight units (5585-X with SSP-60). The Other VPN license allows a maximum of 10,000 IPsec tunnels for one ASA 5585-X with SSP-60. For the entire cluster of eight units, you can only use 10,000 tunnels; the feature does not scale.

Note Traffic for centralized features is forwarded from member units to the control unit over the cluster control link; see "Sizing the Cluster Control Link" in the ASA documentation to ensure adequate bandwidth for the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control units before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control unit.

For centralized features, if the control unit fails, all connections are dropped, and you have to re-establish the connections on the new control unit.

- Site-to-site VPN
- The following application inspections:
  - DCERPC
  - NetBios
  - PPTP
  - RADIUS
  - RSH
  - SUNRPC
  - TFTP
  - XDMCP
- Dynamic routing (spanned EtherChannel mode only)
- Multicast routing (individual interface mode only)
- Static route monitoring
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services

Features Applied to Individual Units

These features are applied to each ASA unit, instead of the cluster as a whole.

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 8 units and with traffic evenly distributed, the conform rate actually becomes 8 times the rate for the cluster.
- Threat detection—Threat detection works on each unit independently; for example, the top statistics is unit-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all units, and one unit will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each unit based on local usage.
- IPS module—There is no configuration sync or state sharing between IPS modules. Some IPS signatures require IPS to keep the state across multiple connections. For example, the port scanning signature is used when the IPS module detects that someone is opening many connections to one server but with different ports. In clustering, those connections will be balanced between multiple ASA devices, each of which has its own IPS module. Because these IPS modules do not share state information, the cluster may not be able to detect port scanning as a result.

Related Topics

- Cluster Information Page, page 3-50

Adding Devices from the Network

One of the easiest and most reliable ways to add devices to the inventory is to identify devices that are active in the network. By providing the IP address (or DNS hostname) of a device, and the credentials required to log into it, Security Manager can obtain much of the information it needs directly from the device, ensuring the accuracy of the information.

Before You Begin

Before beginning this procedure, ensure the following preparations have been made:

- Prepare the devices to be managed by Security Manager. For more information, see Chapter 2, “Preparing Devices for Management”.
- If you are using ACS for authentication, define the devices in ACS. See the *Installation Guide for Cisco Security Manager*.

Related Topics

- Understanding the Device View, page 3-1
- Working with Device Groups, page 3-59
- Viewing or Changing Device Properties, page 3-41

Step 1 In Device view, select **File > New Device** or click the **New Device** button in the Device selector. The New Device wizard opens to the Choose Method page.

Step 2 On the Choose Method page, select **Add Device from Network** and click **Next** to open the Device Information page.

Step 3 On the Device information page, at minimum fill in the following fields. For a detailed explanation of all fields, see Device Information Page – Add Device from Network, page 3-14.

- Enter either a hostname and DNS name, or an IP address (or both).

- Enter a display name, which is the name that will appear in the Security Manager Device selector.
- Select the correct operating system and version. If you are configuring a Catalyst switch or a 7600 router, ensure that you select **IOS - Catalyst Switch/7600** rather than one of the other IOS entries.
- Select the transport protocol that should be used to log into the device, if the device is configured to use a protocol that differs from the default defined in Security Manager. The default is set on the Device Communication administration page (see Device Communication Page, page 11-21).

Click **Next**.

Step 4 On the Device Credentials page, enter the usernames and passwords required to log into the device. Enter at least the primary device credentials, which are the traditional User EXEC mode and Privileged EXEC mode passwords.

For information on the different types of credentials, see Device Credentials Page, page 3-45.

Tip When you click Next or Finished from the Device Credentials page, Security Manager tests whether it can connect to the device. You cannot add the device unless the test succeeds. For more information, see Testing Device Connectivity, page 9-1.

Step 5 (Optional) Click **Next** to open the Device Grouping page, and select the device group to which the imported devices should be added (see Device Groups Page, page 3-49).

Step 6 Click **Finish**. Security Manager opens the Discovery Status dialog box where you can view the status of the device discovery and policy analysis (see Discovery Status Dialog Box, page 5-23).

Note If you are adding a router that is not natively supported and can be generically supported, Security Manager automatically discovers the device as a generic router. You will be prompted whether you want to add the generic device or cancel the operation. For more information about generic device support, see Working with Generically Supported Devices, page 3-9.

Tip If you are discovering policies while adding a device, carefully read any messages that are presented to you. These messages can contain important recommendations on the next steps you should take. For example, when you add Cisco IOS routers or Catalyst devices, we recommend that you immediately deploy the discovered configuration to a file so that Security Manager can take over ownership of the configuration. For more information about deployment methods, see Understanding Deployment Methods, page 8-8.

Step 7 If you are adding a device that contains modules, and Security Manager supports discovering modules for that type of device, you are notified when the discovery of the device chassis is complete and you are asked if you want to discover the device's modules. When you click **Yes**, you are prompted for this information:

- Catalyst 6500 service modules—The Service Module Credentials dialog box opens prompting for the following information, based on the modules contained in the chassis. For more information, see Service Module Credentials Dialog Box, page 3-19.
  - FWSM—The management IP address (recommended), the username and passwords, and the type of discovery you want to perform. If the FWSM is the second device in a failover pair, select **Do Not Discover Module** for the failover module. (Security Manager always manages the active admin context, regardless of whether you added the primary or secondary failover service module.)

  - IDSM—The username and password and the type of discovery you want to perform.
  - ASA-SM—Discovering ASA service modules in a Catalyst 6500 through the chassis is not supported. You must directly add the ASA-SM using the management IP address of the ASA-SM.
- IPS Router Module—The type of discovery you want to perform, the management IP address, the username and password, and other SSL connection information. For more information, see IPS Module Discovery Dialog Box, page 3-20.

You can skip discovery for any module you do not want to manage in Security Manager.

Click **OK**. You are returned to the Discovery Status dialog box, where you can view the progress of service module discovery. When finished, close the window and the device is added to the inventory list. A message will explain if you need to submit the activity for all devices to appear in the list (for example, individual security contexts defined on an ASA device).

Step 8 If you added a device that is managed by an Auto Update Server or Configuration engine, with the device selected in the device selector, select **Tools > Device Properties**. Select the server used with the device in the Auto Update or Configuration Engine settings. You can add the server if it is not listed. For more information, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36.

Device Information Page – Add Device from Network

Use the New Device wizard's Device Information page for adding devices from the network to specify the device's identifying information.

Navigation Path

To start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.

Related Topics

- Understanding the Device View, page 3-1
- Adding Devices from the Network, page 3-12
- Device Credentials Page, page 3-45
- Device Groups Page, page 3-49
- Discovering Policies, page 5-12
- Device Communication Page, page 11-21

Field Reference

Table 3-2 New Device Wizard, Device Information Page When Adding Devices from the Network

Element	Description
Identity	
IP Type	Whether the IP address for the device is static (defined on the device) or dynamic (supplied by a DHCP server). Depending on the IP type you select, the displayed fields differ. You can add only devices that have static IP addresses. If you want to add a device that uses dynamic addresses (supplied by a DHCP server), determine the current IP address for the device, use that address, and after adding the device, update its properties to change the IP Type to Dynamic and to identify the AUS or Configuration Engine that is managing the device. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. The IPv6 address is a 128-bit unique address. For IPv6 address, only Static IP Type is supported. Dynamic IP Type is not supported for IPv6 addresses.
Hostname	The DNS hostname for the device. Enter the DNS hostname if the IP address is not known. **Note** You must enter either the DNS hostname or the IP address, or both.
Domain Name	The DNS domain name for the device.
IP Address	The management IP address of the device. The IP address must be in the dotted quad format, for example, 10.64.3.8. **Note** You must enter either the IP address or the DNS hostname, or both. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. If a device is configured in dual stack, Security Manager would communicate with the device based on the device's IP address added in Security Manager. The IPv6 address is a 128-bit unique address.
Display Name	The name to display in the Security Manager Device selector. If you enter a hostname or IP address, it is entered automatically in this field, but you can change it. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; and the following characters: _ -. : and space. **Note** Two devices cannot have the same display name.

Table 3-2 New Device Wizard, Device Information Page When Adding Devices from the Network (continued)

Element	Description
OS Type	The family of the operating system running on the device. You must be careful to select the correct type, because your selection affects how Security Manager tries to log into the device and obtain its configuration. The options are: • **IOS 12.3+**—For Cisco routers running Cisco IOS Software Release 12.3 or higher. Do not select this for Catalyst 6500/7600 or other Catalyst devices. **Tip** Select this option for Aggregation Services Routers (ASR) even if they are running a version of 12.2. The ASR IOS releases are treated as higher releases. • **IOS - 12.2, 12.1**—For Cisco routers running Cisco IOS Software Releases 12.2 or 12.1. Do not select this for ASRs, Catalyst 6500/7600, or other Catalyst devices. • **IOS - Catalyst Switch/7600**—For all Catalyst switches and 7600 devices. • **ASA**—For all ASA devices. • **FWSM**—For all FWSM devices. • **IPS**—For all devices running the IPS software. • **PIX**—For all PIX devices. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. This feature is available only for devices where the Operating System type is ASA or FWSM.
Transport Protocol	The protocol Security Manager should use when connecting to the device. Select a protocol that is configured on the device and for which you can supply credentials. Each device type has a default protocol that is the method normally used with the device.

Table 3-2 New Device Wizard, Device Information Page When Adding Devices from the Network (continued)

Element	Description
System Context	Whether to discover the system execution space of a PIX Firewall 7, ASA, or FWSM device that is running in multiple-context mode. If you are discovering a device that hosts multiple security contexts, whether you select this checkbox has important implications in how you can configure the device in Security Manager. What gets discovered on the device also depends on whether you select the **Discover Policies for Security Contexts** checkbox. • Both **System Context** and **Discover Policies for Security Contexts** selected—This is the recommended selection. Security Manager discovers the system execution space and all of the security contexts defined on the device, and lists them in the device selector. The base display name represents the system execution space (for example, 10.10.11.24), whereas the security contexts are represented by nodes with the context name appended to the device name (for example, 10.10.11.24_admin), unless you changed the default naming convention configured on the Discovery page (see Discovery Page, page 11-25). • **System Context** selected, **Discover Policies for Security Contexts** deselected—The system execution space is discovered and added to the device selector. You can then discover the policies for the security contexts at a later time. This method might be appropriate if you have one group of people who discover inventory and another group that discovers policies. • Neither checkbox selected—Only the Admin context gets discovered and added to the device selector. You cannot discover the other security contexts or manage them.

Table 3-2 New Device Wizard, Device Information Page When Adding Devices from the Network (continued)

Element	Description
Discover Device Settings	
Discover	The type of elements that should be discovered and added to the inventory. You have these options: • Policies and Inventory—Discover policies, interfaces, and service modules (if applicable). This is the default and recommended option. When policy discovery is initiated, the system analyzes the configuration on the device, then imports the configured service and platform policies. When inventory discovery is initiated, the system analyzes the interfaces on the device and then imports the interface list. If the device is a composite device, all the service modules in the device are discovered and imported. If you select this option, the checkboxes below are activated and you can use them to control the types of policies that are discovered. **Note** During discovery, if you import an ACL that is inactive, it is shown as disabled in Security Manager. If you deploy the same ACL, it will be removed by Security Manager. • Inventory Only—Discovers interfaces and service modules (if applicable). • No Discovery—All discovery is skipped. No policy, interface, or service module information for the device is added to the device inventory.
Platform Settings	Whether to discover the platform settings, which are also called platform-specific policy domains. Platform-specific policy domains exist on firewall devices and Cisco IOS routers. These domains contain policies that configure features that are specific to the selected platform. For more information, see Service Policies vs. Platform-Specific Policies, page 5-2.
Firewall Policies	Whether to discover firewall policies, which are also called firewall services. Firewall services include policies such as access rules, inspection rules, AAA rules, web filter rules, and transparent rules. For details see, Chapter 12, “Introduction to Firewall Services”
IPS Policies	Whether to discover IPS policies such as signatures and virtual sensors. For more information, see Overview of IPS Configuration, page 36-5 or Overview of Cisco IOS IPS Configuration, page 45-4.
RA VPN Policies	Whether to discover IPSec and SSL remote access VPN policies such as IKE proposals and IPsec proposals. This option is disabled if the device does not support remote access VPN configuration. For more information, see Chapter 30, “Managing Remote Access VPNs: The Basics”.

Table 3-2 New Device Wizard, Device Information Page When Adding Devices from the Network (continued)

Element	Description
Discover Policies for Security Context	Whether to discover policies for security contexts. Security contexts apply to PIX Firewall, ASA, or FWSM devices. This field is active only if you select **Static** for IP Type and **System Context**.

Service Module Credentials Dialog Box

Use the Service Module Credentials dialog box to add the credentials required to log into supported service modules in a Catalyst device.

The dialog box includes a group for each slot that contains a supported module, and the type of module is indicated. For example, a group might be called **Slot 3 (IDSM) Credentials**, which indicates that there is an IDSM in the third slot of the chassis.

Note Although Security Manager discovers VPN modules, the discovery is done through the chassis and no credentials are required. ASA service modules (ASA-SM) cannot be discovered through the chassis; you must add them individually.

Navigation Path

After you discover policies on a Catalyst chassis that can contain service modules, you are asked if you want to discover its service modules. If you click **Yes**, this dialog box appears. You can perform policy discovery using any of these methods:

- When adding a device from the network. See Adding Devices from the Network, page 3-12.
- When adding devices from an export file. See Adding Devices from an Inventory File, page 3-31.
- When performing policy discovery on a device that is already in the inventory. See Discovering Policies on Devices Already in Security Manager, page 5-15.

Related Topics

- Chapter 59, “Configuring Security Contexts on Firewall Devices”

Field Reference

Table 3-3 Service Module Credentials Dialog Box

Element	Description
Discovery Mode	The types of policies to discovery for this module: • Discover Inventory and Policies—Discover inventory and security policies. This is the recommended option. • Discover Inventory Only—Do not discover security policies, but discover inventory, such as VLAN configuration, security contexts, and interfaces. You can discover the policy configuration later by right-clicking the service module and then selecting **Discover Policies on Device**. • Do Not Discover Module—Skip discovery on this module and do not add it to the inventory.

Table 3-3 Service Module Credentials Dialog Box (continued)

Element	Description
Connect to FWSM	How Security Manager should access the FWSM: • Directly—Connect to the FWSM using its management IP address. This is the recommended approach. It is the required method if you are connecting to a failover device; otherwise, Security Manager might connect to a standby FWSM after a failover. • via Chassis—Connect to the FWSM through the chassis. This method has the restriction that there should be fewer than 20 security contexts defined on the FWSM. Security Manager connects to the Catalyst device through SSH and then to the FWSM through the **session** command. The number of concurrent SSH sessions is limited on a Catalyst device, with a default of 5. Policy discovery uses one SSH session for each security context, so a large number of contexts might lead to connection failures. If you select **Directly**, Security Manager connects to the FWSM through SSL, which has a greater concurrent session limit.
Management IP	The management IP address for the service module. For FWSMs, this field is not available if you select **via Chassis** for the connection method.
Username	The user name for the service module. For FWSMs running in multiple-context mode, a footnote explains which context's username and password to enter, either the system or the admin context. If you are connecting to a multiple-context mode device through the switch chassis, you must configure the same username and password for both the system execution space and the admin context, and specify those credentials in this dialog box. User names be at least four characters. Passwords can be three to 32 characters; we recommend that passwords be at least eight characters. For ASA devices running the software version 9.6(1) or later, you can enter a password up to 127 characters.
Password	The User EXEC mode password for the service module. In the Confirm field, enter the password again.
Enable Password (FWSM only)	The Privileged EXEC mode password for the service module. In the Confirm field, enter the password again.

IPS Module Discovery Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

Use the IPS Module Discovery dialog box to add the credentials required to log into an IPS module, such as an AIM-IPS or NME, on a router you are adding to the inventory.

Navigation Path

After you discover policies on a router chassis that contains an IPS module, you are asked if you want to discover its modules. If you click **Yes**, this dialog box appears. You can perform policy discovery using any of these methods:

- When adding a device from the network. See Adding Devices from the Network, page 3-12.
- When adding devices from an inventory file. See Adding Devices from an Inventory File, page 3-31.
- When performing policy discovery on a device that is already in the network. See Discovering Policies on Devices Already in Security Manager, page 5-15.

Field Reference

Table 3-4 IPS Module Discovery Dialog Box

Element	Description
Discovery	The type of discovery for this module: • Discover Inventory and Policies—Discover inventory and security policies. This is the recommended option. • Discover Inventory Only—Do not discover security policies, but discover inventory, such as virtual sensors and interfaces. You can discover the policy configuration later by right-clicking the module and selecting **Discover Policies on Device**. • Do Not Discover Module—Skip discovery on this module and do not add it to the inventory.
IP Address	The management IP address for the module.
HTTP Credentials Group The credentials required to log into the module.	
Username	The username for the module.
Password	The password for the specified username. In the Confirm field, enter the password again.
HTTP Port	The port configured for HTTP access to the module. The default is 80.
HTTPS Port	The port configured for SSL (HTTPS) access to the module. The default is defined on the Device Communication page (**Tools > Security Manager Administration > Device Communication**, for more information, see Device Communication Page, page 11-21). The port typically used is 443. To override the default, deselect **Use Default** and enter the correct port number.
IPS RDEP Mode	The connection method to use for contacting IPS devices when making RDEP or SDEE connections (for event monitoring).
Certificate Common Name	The name assigned to the certificate. The common name can be the name of a person, system, or other entity that was assigned to the certificate. In the Confirm field, enter the common name again.

Adding Devices from Configuration Files

You can add devices to the inventory by having Security Manager process the device configurations without logging into the devices. For each device, you must copy the device configuration to a file and put the file on the Security Manager server.

You cannot use this procedure to add IPS or Catalyst 6500/7600 devices to the inventory.

Before You Begin

Before beginning this procedure, ensure the following preparations have been made:

- Prepare the devices to be managed by Security Manager. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are using ACS for authentication, define the devices in ACS. See the *Installation Guide for Cisco Security Manager*.
- Copy the device configuration files to a directory on the Security Manager server. You cannot use a mounted drive. Use a naming convention that will help you select the correct device type for each configuration.

Related Topics

- Understanding the Device View, page 3-1
- Working with Device Groups, page 3-59
- Viewing or Changing Device Properties, page 3-41

Step 1 In Device view, select **File > New Device** or click the **New Device** button in the Device selector. The New Device wizard opens to the Choose Method page.

Step 2 On the Choose Method page, select **Add from Configuration File** and click **Next** to open the Device Information page (see Device Information Page—Configuration File, page 3-23).

Step 3 Select the device type for the configuration files from the Device Type selector, and select the appropriate system object ID. If you have configuration files for more than one device type, add them in batches based on device type.

Note If you are adding a router that does not appear in the explicitly supported device list, you can try to manage it as a generic device using the following device types: Cisco Generic Aggregation Services Router—for devices running Cisco IOS XE Software and Cisco Generic Integrated Services Router—for devices running Cisco IOS Software. For more information about generic device support, see Working with Generically Supported Devices, page 3-9.

Step 4 Click **Browse** and select the configuration files that contain the devices (of the specified type) that you want to add.

Step 5 Select the appropriate discovery options to indicate which types of policies you want to discover, if any.

Step 6 (Optional) Click **Next** and select the device groups to which the new devices should belong.

Step 7 Click **Finish**. Security Manager opens the Discovery Status dialog box where you can view the status of the configuration file analysis (see Discovery Status Dialog Box, page 5-23). When finished, close the window and the device is added to the inventory list.

Tip If you are discovering policies and get unexpected errors, it might be because the configuration file includes only the major Cisco IOS software version and not the point release information. Some policies defined on the device might use features that became available in a point release, which means that Security Manager might not recognize them as being supported. To resolve the problem, after adding the device, select it in the Device selector, right-click, and select **Device Properties**. On the General page, update the **Target OS Version** field with the software version closest to the one running on the device without being higher than it (you can get the version number using the **show version** command on the device's CLI). You can then rediscover policies by right-clicking and selecting **Discover Policies on Device**.

Step 8 If you added a device that is managed by an Auto Update Server or Configuration engine, with the device selected in the device selector, select **Tools > Device Properties**. Select the server used with the device in the Auto Update or Configuration Engine settings. You can add the server if it is not listed. For more information, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36.

Device Information Page—Configuration File

Use the New Device wizard's Device Information page for adding devices from configuration files to select the configuration files and to specify policy discovery options.

Navigation Path

To start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.

Related Topics

- Understanding the Device View, page 3-1
- Adding Devices from Configuration Files, page 3-22
- Device Groups Page, page 3-49
- Discovering Policies, page 5-12
- Discovery Status Dialog Box, page 5-23

Field Reference

Table 3-5 New Device Wizard, Device Information Page When Adding Devices from Configuration Files

Element	Description
Device Type selector	Organizes the devices by device-type and device-family. Select the device type for the new device. You must select the correct device type for the configuration file you are adding. **Note** If you are adding a router that does not appear in the explicitly supported device list, you can try to manage it as a generic device using the following device types: Cisco Generic Aggregation Services Router—for devices running Cisco IOS XE Software and Cisco Generic Integrated Services Router—for devices running Cisco IOS Software. For more information about generic device support, see Working with Generically Supported Devices, page 3-9.
System Object ID	The system object identifiers for the device type you selected from the Device Type selector. Select the correct ID for your device.
Configuration Files	The configuration files from the devices you are adding to the inventory. You can specify more than one configuration file, but they must all be for the same device type. Separate the file names with commas. For ASA, PIX, and FWSM devices that have multiple security contexts, keep in mind that there are separate configuration files for each security context and the system execution space (the system context). Select the configuration file for the system execution space to add the base device. Click **Browse** to select the files from the Security Manager server, or manually type in the file names (including the full path). For information on selecting files, see Selecting or Specifying a File or Directory in Security Manager, page 1-50.
Options	The additional options available on the device. Select IPS if the IPS feature is available on the device.
Target OS Version (Routers only.)	The OS version on which you want to base the router's configuration. • Select Default to read the OS version from the configuration file. • Select a specific release number if you want to ensure the desired release is used. If the desired release is not listed, select the release number that is closet but lower than the release running on the device. You might have to select a specific release when discovering configuration files for Aggregation Services Routers (ASR) or for configuration files that include zone-based firewall policies.
License Supports Failover (ASA 5505, 5510 only.)	Whether an optional failover license is installed on the device. The option is active for ASA 5505 and 5510 devices only. Security Manager deploys failover policies to the device only if this option is selected. **Tip** If you discover policies from the device, Security Manager determines the license status and sets this option appropriately.

Table 3-5 New Device Wizard, Device Information Page When Adding Devices from Configuration Files (continued)

Element	Description
Discover Device Settings	
Discover	The type of elements that should be discovered and added to the inventory. You have these options: • Policies and Inventory—Discover policies, interfaces, and service modules (if applicable). This is the default and recommended option. When policy discovery is initiated, the system analyzes the configuration file, then imports the configured service and platform policies. When inventory discovery is initiated, the system analyzes the interfaces defined in the file and then imports the interface list. If you select this option, the checkboxes below are activated and you can use them to control the types of policies that are discovered. **Note** During discovery, if you import an ACL that is inactive, it is shown as disabled in Security Manager. If you deploy the same ACL, it will be removed by Security Manager. • Inventory Only—Discovers interfaces and service modules (if applicable). • No Discovery—All discovery is skipped. No policy, interface, or service module information for the device is added to the device inventory.
Platform Settings	Whether to discover the platform settings, which are also called platform-specific policy domains. Platform-specific policy domains exist on firewall devices and Cisco IOS routers. These domains contain policies that configure features that are specific to the selected platform. For more information, see Service Policies vs. Platform-Specific Policies, page 5-2.
Firewall Policies	Whether to discover firewall policies, which are also called firewall services. Firewall services include policies such as access rules, inspection rules, AAA rules, web filter rules, and transparent rules. For details see, Chapter 12, “Introduction to Firewall Services”
IPS Policies	Whether to discover IPS policies such as signatures and virtual sensors. For more information, see Overview of IPS Configuration, page 36-5 or Overview of Cisco IOS IPS Configuration, page 45-4.
RA VPN Policies	Whether to discover IPSec and SSL remote access VPN policies such as IKE proposals and IPsec proposals. This option is disabled if the device does not support remote access VPN configuration. For more information, see Chapter 30, “Managing Remote Access VPNs: The Basics”.

Adding Devices by Manual Definition

If a device is not yet active on the network, you can add it to Security Manager and preprovision a configuration for the device. In general, you should not use manual definition for a device that exists in the network, because it is much easier to use one of the other techniques for adding devices.

Before You Begin

Before beginning this procedure, ensure the following preparations have been made:

- Prepare the devices to be managed by Security Manager. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are using ACS for authentication, define the devices in ACS. See the *Installation Guide for Cisco Security Manager*.

Related Topics

- Understanding the Device View, page 3-1
- Working with Device Groups, page 3-59
- Viewing or Changing Device Properties, page 3-41

Step 1 In Device view, select **File > New Device** or click the **New Device** button in the Device selector. The New Device wizard opens to the Choose Method page.

Step 2 On the Choose Method page, select **Add New Device** and click **Next** to open the Device Information page.

Step 3 On the Device Information page, at minimum fill in the following fields. For a detailed explanation of all fields, see Device Information Page—New Device, page 3-27.

- Select the device type from the Device Type selector at the left of the page, and select the system object ID at the bottom of the selector.

  Note If you are adding a router that does not appear in the explicitly supported device list, you can try to manage it as a generic device using the following device types: Cisco Generic Aggregation Services Router—for devices running Cisco IOS XE Software and Cisco Generic Integrated Services Router—for devices running Cisco IOS Software. For more information about generic device support, see Working with Generically Supported Devices, page 3-9.

- In the IP Type field, select whether the device uses a static address (the IP address is defined on the device) or a dynamic one (the address is provided by a DHCP server).
- For devices with static addresses, enter either a DNS hostname and domain name, or an IP address (or both).
- Enter a display name, which is the name that will appear in the Security Manager Device selector.
- Ensure that the correct operating system and version are selected.
- If you use a server to manage configurations for the device, which is required for dynamically addressed devices, select the Auto Update Server or Configuration Engine that manages the device and enter the device identity string the server uses for the device. If the server is not listed, select **Add Server** and add it to the inventory. For information on adding servers, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36

When you are finished filling in the device information, click **Next** to proceed to the Device Credentials page.

Step 4 (Optional) On the Device Credentials page, enter the usernames and passwords required to log into the device. Typically, you need to enter the primary device credentials, which are the traditional User EXEC mode and Privileged EXEC mode passwords. If you do not enter credentials, you can add them later on the Device Properties page.

For information on the different types of credentials, see Device Credentials Page, page 3-45.

Click **Next**.

Step 5 (Optional) On the Device Grouping page, select the group to which the device should belong, if any. See Device Groups Page, page 3-49

Step 6 Click **Finish**. The device is added to the inventory.

Tip If you are adding a PIX, ASA, or FWSM device, you should discover the factory default settings for the device and its security contexts. For more information, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Device Information Page—New Device

Use the New Device wizard's Device Information page for adding new devices (that do not yet exist in the network) to specify the device's identifying information.

Navigation Path

To start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.

Related Topics

- Understanding the Device View, page 3-1
- Adding Devices by Manual Definition, page 3-26
- Device Credentials Page, page 3-45
- Device Groups Page, page 3-49

Field Reference

Table 3-6 New Device Wizard, Device Information Page When Adding New Devices

Element	Description
Device Type	
Device Type selector	Organizes the devices by device-type and device-family. Select the device type for the new device. **Note** If you are adding a router that does not appear in the explicitly supported device list, you can try to manage it as a generic device using the following device types: Cisco Generic Aggregation Services Router—for devices running Cisco IOS XE Software and Cisco Generic Integrated Services Router—for devices running Cisco IOS Software. For more information about generic device support, see Working with Generically Supported Devices, page 3-9.
System Object ID	The system object identifiers for the device type you selected from the Device Type selector. Select the correct ID for your device.
Identity	
IP Type	Whether the IP address for the device is static (defined on the device) or dynamic (supplied by a DHCP server). Depending on the IP type you select, the displayed fields differ. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. The IPv6 address is a 128-bit unique address. For IPv6 address, only Static IP Type is supported. Dynamic IP Type is not supported for IPv6 addresses.
Hostname (Static IP only)	The DNS hostname for the device. Enter the DNS hostname if the IP address is not known. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; and hyphen (-). **Note** You must enter either the DNS hostname or the IP address, or both. Two devices cannot have the same DNS hostname and domain name combination.
Domain Name (Static IP only)	The DNS domain name for the device. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; period (.) and hyphen (-).

Table 3-6 New Device Wizard, Device Information Page When Adding New Devices (continued)

Element	Description
IP Address (Static IP only)	The management IP address of the device. The IP address must be in the dotted quad format, for example 10.64.3.8. **Note** You must enter either the IP address or the DNS hostname, or both. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. If a device is configured in dual stack, Security Manager would communicate with the device based on the device's IP address added in Security Manager. The IPv6 address is a 128-bit unique address.
Display Name	The name to display in the Security Manager Device selector. If you enter a hostname or IP address, it is entered automatically in this field, but you can change it. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; and the following characters: _ -. : and space. **Note** Two devices cannot have the same display name.
Operating System	
OS Type	The type of operating system. Based on the device type, the OS type is selected automatically. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. This feature is available only for devices where the Operating System type is ASA or FWSM.
Image Name	The name of the image that will run on the device.
Target OS Version	The target OS version for which you want to apply the configuration. This selection determines the type of commands used when Security Manager generates configuration files.
Options	The additional options available on the device. Select IPS if the IPS feature is available on the device.
Contexts	Whether the device hosts a single security context (Single) or multiple security contexts (Multi). This field is displayed only if the OS type is an FWSM, ASA, or PIX Firewall 7.0.
Operational Mode	The mode in which the device is operating. This field is displayed only if the OS type is FWSM, ASA, or PIX Firewall 7.0+. The options available are: Transparent or Router. If you choose Multi for Contexts, this mode defaults to Mixed. Mixed applies only to ASA 9.0+ and FWSM 3.1+ devices, and ASA-SMs.

Table 3-6 New Device Wizard, Device Information Page When Adding New Devices (continued)

Element	Description
Auto Update or Configuration Engine	
This group is named differently depending on the device type you select: • Auto Update—For PIX Firewall and ASA devices. • Configuration Engine—For Cisco IOS Routers. Use these fields to identify the server that manages the device, if any. A server is required for a device with a dynamic IP address. You cannot define a server for Catalyst 6500/7600 or FWSM devices.	
Server	The Auto Update Server or Configuration Engine that manages the device. You can add servers to the list by selecting **Add Servers**, which opens the Server Properties dialog box (see Server Properties Dialog Box, page 3-38). You can also edit the properties of a server by selecting **Edit Server**, which opens the Available Servers dialog box (see Available Servers Dialog Box, page 3-39). For more information on managing this list of servers, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36
Device Identity	The string value that uniquely identifies the device in Auto Update Server or the Configuration Engine.
Additional Fields	
Manage in Cisco Security Manager	Whether Security Manager manages the device. This check box is selected by default. If the only function of the device you are adding is to serve as a VPN end point, deselect this check box. Security Manager will not manage configurations nor will it upload or download configurations on this device. For more information, see Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11.
Security Context of Unmanaged Device	Whether to manage a security context whose parent (the PIX Firewall, ASA, or FWSM device) is not managed by Security Manager. This field is active only if the device you selected in the Device selector is a firewall device, such as PIX Firewall, ASA, or FWSM and that firewall device supports security contexts. You can partition a PIX Firewall, ASA, or FWSM into multiple security firewalls, also known as security contexts. Each context is an independent system with its own configuration and policies. You can manage these standalone contexts in Security Manager, even though the parent device is not managed by Security Manager. For more information, see Chapter 59, "Configuring Security Contexts on Firewall Devices". **Note** If you select this check box, the available target OS version for the security module is displayed in the Target OS Version field.

Table 3-6 New Device Wizard, Device Information Page When Adding New Devices (continued)

Element	Description
License Supports Failover (ASA 5505, 5510 only.)	Whether an optional failover license is installed on the device. The option is active for ASA 5505 and 5510 devices only. Security Manager deploys failover policies to the device only if this option is selected. **Tip** If you discover policies from the device, Security Manager determines the license status and sets this option appropriately.

Adding Devices from an Inventory File

You can add devices from an inventory file in comma-separated values (CSV) format. For example, an inventory file you exported from CiscoWorks Common Services Device Credential Repository (DCR) or another Security Manager server, or the seed file you used with Cisco Security Monitoring, Analysis and Response System (CS-MARS). For detailed information about the inventory file formats, see Supported CSV Formats for Inventory Import/Export, page 10-9.

Tips

- This procedure explains how to use a CSV file for importing devices. If you have a .dev file, which includes not only the inventory but the policies and policy objects assigned to the devices, you cannot use this procedure. Instead, use the **File > Import** command and follow the instructions in Importing Policies or Devices, page 10-13.
- If you want to build an inventory file by hand, the easiest approach is to export the Security Manager inventory in the desired format and use that file as the basis for your inventory file.
- The devices you import cannot be duplicates of devices already in the device inventory. You cannot, for example, update device information in the inventory by re-importing the device.
- If you are adding a router that does not appear in the explicitly supported device list, you can try to manage it as a generic device. When adding a generic router from an inventory file, you must make sure that you use the correct SYSOBJID in the CSV file. Use 1.0.0.0.0.0.0.1 for devices running Cisco IOS Software (Cisco Generic Integrated Services Router) and use 1.0.0.0.0.0.0.2 for devices running Cisco IOS XE Software (Cisco Generic Aggregation Services Router). For more information about generic device support, see Working with Generically Supported Devices, page 3-9.

Before You Begin

Before beginning this procedure, ensure the following preparations have been made:

- Prepare the devices to be managed by Security Manager. For more information, see Chapter 2, “Preparing Devices for Management”.
- If you are using ACS for authentication, define the devices in ACS. See the *Installation Guide for Cisco Security Manager*.
- Put the inventory file you want to use on the Security Manager server. You cannot import devices from a file on your client system.
- If you are using a non-standard communication protocol for a type of device, update the global device communication properties to specify the correct protocol. For more information, see Device Communication Page, page 11-21.

Related Topics

- Understanding the Device View, page 3-1
- Working with Device Groups, page 3-59
- Viewing or Changing Device Properties, page 3-41

Step 1 In Device view, select **File > New Device** or click the **New Device** button in the Device selector. The New Device wizard opens to the Choose Method page.

Step 2 On the Choose Method page, select **Add Device from File** and click **Next** to open the Device Information page (see Device Information Page—Add Device from File, page 3-33).

Step 3 Click **Browse** and select the inventory file that contains the devices that you want to import. Make sure that you select the correct file type to indicate how the file is formatted.

Security Manager evaluates the contents of the inventory file and displays the list of devices in the import table. All devices that have the status Ready to Import are automatically selected. The list identifies the reasons the unselected devices cannot be imported. You can deselect any devices that you do not want to import.

To see detailed information on a device, select it in the import table. The details are displayed in the bottom pane. You can select different discovery options or transport settings per device.

Tip If you selected an inventory file in the Security Manager format, you have the option to import the devices without performing policy discovery. This makes it possible for you to add devices that are not currently active in the network. If you want to perform policy discovery on a device, select the device, select **Perform Device Discovery** in the bottom panel, and select your discovery options. You can select policy discovery settings for all devices in a folder by selecting the folder instead of individual devices. The other CSV formats require that you perform policy discovery during import.

When you are finished analyzing the list and modifying discovery and transport settings, click **Next** to continue to the optional step of selecting groups, or click **Finish** to complete the wizard. In either case, Security Manager attempts to log into each device and perform the discovery you selected unless you are using a CSV file in Security Manager format and elected not to perform discovery. For the other formats, Security Manager must be able to log into the device to add it to the inventory. The status is displayed in the Discovery Status dialog box (see Discovery Status Dialog Box, page 5-23).

Tip If you are discovering policies while adding a device, carefully read any messages that are presented. These messages can contain important recommendations on the next steps you should take. For example, when you add Cisco IOS routers or Catalyst devices, we recommend that you immediately deploy the discovered configuration to a file so that Security Manager can take ownership of the configuration. For more information about deployment methods, see Understanding Deployment Methods, page 8-8.

Step 4 (Optional) On the Device Grouping page, select the device group to which the imported devices should be added (see Device Groups Page, page 3-49).

Click **Finish**.

Step 5 If you are adding a device that contains modules and you are performing device discovery, and Security Manager supports discovering modules for that type of device, you are notified when the discovery of the device chassis is complete and you are asked if you want to discover the device's modules. When you click **Yes**, you are prompted for this information:

- Catalyst 6500 service modules—The Service Module Credentials dialog box opens prompting for the following information, based on the modules contained in the chassis. For more information, see Service Module Credentials Dialog Box, page 3-19.
  - FWSM—The management IP address (recommended), the user name and passwords, and the type of discovery you want to perform. If the FWSM is the second device in a failover pair, select **Do Not Discover Module** for the failover module. (Security Manager always manages the active admin context, regardless of whether you added the primary or secondary failover service module.)
  - IDSM—The user name and password and the type of discovery you want to perform.
  - ASA-SM—Discovering ASA service modules in a Catalyst 6500 through the chassis is not supported. You must directly add the ASA-SM using the management IP address of the ASA-SM.
- IPS Router Module—The type of discovery you want to perform, the management IP address, the user name and password, and other SSL connection information. For more information, see IPS Module Discovery Dialog Box, page 3-20.

You can skip discovery for any module you do not want to manage in Security Manager.

Click **OK**. You are returned to the Discovery Status dialog box, where you can view the progress of service module discovery.

Device Information Page—Add Device from File

Use the New Device wizard's Device Information page for adding devices from an inventory file to select the file and to specify policy discovery options. The inventory file must be on the Security Manager server; you cannot use an inventory file on a client system.

The formats you can use for the inventory file are explained in Supported CSV Formats for Inventory Import/Export, page 10-9. Typically, the inventory file will have been exported from another Security Manager server, from a CiscoWorks Common Services server, or it will be the seed file used to populate the inventory of a Cisco Security Monitoring, Analysis and Response System (CS-MARS) server.

If you are trying to import devices using a .dev file, you need to use the File > Import command instead of this page. For more information, see Importing Policies or Devices, page 10-13.

Tip If you are adding devices that contain modules, for example, a Catalyst switch with an FWSM, you are prompted for module discovery information after you click **Finish**.

Navigation Path

To start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.

Related Topics

- Understanding the Device View, page 3-1
- Adding Devices from an Inventory File, page 3-31
- Device Groups Page, page 3-49
- Discovering Policies, page 5-12
- Device Communication Page, page 11-21

- Discovery Status Dialog Box, page 5-23

Field Reference

Table 3-7 New Device Wizard, Device Information Page When Adding Devices from Inventory Files

Element	Description
Import Devices From	The inventory file that contains the devices you want to import. Click **Browse** to select the file on the Security Manager server. When selecting the file, you must also select the correct file type so that Security Manager can correctly evaluate the comma-separated values (CSV) file.
Device Import Table	
After you select a file, Security Manager evaluates its contents and displays the list of devices defined in the file in the table in the upper pane of the page. Security Manager automatically selects all devices whose status is Ready to Import. Typically, these are the devices that do not already exist in the device inventory. The table contains the following columns.	
Import	Select this checkbox to add the device to the inventory. You can select or deselect a folder to select or deselect all devices within the folder.
Display Name	The name that will be displayed in the Security Manager Device selector.
Host Name	The host name defined on the device.
Transport	The transport protocol that should be used to connect to the device.
Status	Whether Security Manager can import the device. Devices can be imported only if they have the status Ready to Import. For detailed information on a device's status, select it and read the expanded status information in the Status text box in the lower right corner of the page.
Device Type	The type of device.
Details Pane	
Below the device import table is a pane that displays the details for the device selected in the table. The Identity information repeats the table fields. The Status text box displays an extended explanation of the import status. The Discover Device Settings and Transport groups let you specify how Security Manager should import the device. If you select a folder instead of a device, the settings you select apply to all devices in the folder. The settings are explained below.	
Discover Device Settings	
Perform Device Discovery	Whether to discover policies directly from the device: • If the inventory file is in Security Manager format, you must select Perform Device Discovery to discovery inventory and policies (otherwise, the device is added without being evaluated). If you are adding offline or standby devices, you can leave this option deselected to easily add the device to the inventory. • All other inventory file types require device discovery.

Table 3-7 New Device Wizard, Device Information Page When Adding Devices from Inventory Files (continued)

Element	Description
System Context	Whether the selected device is the system execution space on a device running in multiple context mode (that is, more than one security context is defined on the device). If the device is the system execution space, you must select this option for discovery to complete correctly.
Discover	The type of elements that should be discovered and added to the inventory. You have these options: • Policies and Inventory—Discover policies, interfaces, and service modules (if applicable). This is the default and recommended option. When policy discovery is initiated, the system analyzes the configuration on the device, then imports the configured service and platform policies. When inventory discovery is initiated, the system analyzes the interfaces on the device and then imports the interface list. If the device is a composite device, all the service modules in the device are discovered and imported. If you select this option, the checkboxes below are activated and you can use them to control the types of policies that are discovered. **Note** During discovery, if you import an ACL that is inactive, it is shown as disabled in Security Manager. If you deploy the same ACL, it will be removed by Security Manager. • Inventory Only—Discovers interfaces and service modules (if applicable).
Platform Settings	Whether to discover the platform settings, which are also called platform-specific policy domains. Platform-specific policy domains exist on firewall devices and Cisco IOS routers. These domains contain policies that configure features that are specific to the selected platform. For more information, see Service Policies vs. Platform-Specific Policies, page 5-2.
Firewall Policies	Whether to discover firewall policies, which are also called firewall services. Firewall services include policies such as access rules, inspection rules, AAA rules, web filter rules, and transparent rules. For details see, Chapter 12, “Introduction to Firewall Services”
IPS Policies	Whether to discover IPS policies such as signatures and virtual sensors. For more information, see Overview of IPS Configuration, page 36-5 or Overview of Cisco IOS IPS Configuration, page 45-4.
RA VPN Policies	Whether to discover IPSec and SSL remote access VPN policies such as IKE proposals and IPsec proposals. This option is disabled if the device does not support remote access VPN configuration. For more information, see Chapter 30, “Managing Remote Access VPNs: The Basics”.
Discover Policies for Security Contexts	For devices running in multiple-context mode, where more than one security context is defined on the device, whether to discover those security contexts.

Table 3-7 New Device Wizard, Device Information Page When Adding Devices from Inventory Files (continued)

Element	Description
Transport	
The transport settings determine the method Security Manager will use to contact the device. Each device type has a default method, but you can select your preferred transport method. The device must be configured to respond to the method you select. If you are not performing device discovery, the device is not contacted.	
Protocol	The protocol Security Manager should use when connecting to the device.
Server	For devices that use them, the name of the Auto Update Server (AUS) or Configuration Engine server the device uses to obtain configuration updates. The server must already be defined in Security Manager, or you must select the server from the import list, to import devices that use these servers.
Device Identity	For devices that use servers, the string value that uniquely identifies the device in the Auto Update Server or the Configuration Engine.

Working with the Device Inventory

The following topics describe tasks related to managing the device inventory.

- Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36
- Adding or Changing Interface Modules, page 3-40
- Viewing or Changing Device Properties, page 3-41
- Changing Critical Device Properties, page 3-52
- Showing Device Containment, page 3-56
- Cloning a Device, page 3-56
- Deleting Devices from the Security Manager Inventory, page 3-58

In addition to these topics, see the following related topics:.

- Adding Devices to the Device Inventory, page 3-7
- Exporting the Device Inventory, page 10-6
- Importing Policies or Devices, page 10-13

Adding, Editing, or Deleting Auto Update Servers or Configuration Engines

If you want to use Security Manager to manage devices that use other servers to manage their configuration (for example, devices that have dynamic IP addresses supplied by a DHCP server, an address that might not stay constant between device reboots), you must identify the server in Security Manager. These are the servers you can use:

- Auto Update Server (AUS), which is used for upgrading device configuration files on PIX Firewall and ASA devices that use the auto update feature.

- Cisco Configuration Engine, which is used for upgrading device configuration files on Cisco IOS routers, ASA devices, and PIX Firewalls that use the configuration engine feature.

Security Manager cannot initiate direct communication with devices that acquire their interface addresses using DHCP because their IP addresses are not known ahead of time. Furthermore, these devices might not be running, or they might be behind firewalls and NAT boundaries when the management system must make changes. These devices connect to an Auto Update Server or Configuration Engine to get device information.

You can add AUS and Configuration Engine servers to the device inventory when you add devices manually or when you view device properties. You do not have to be adding or viewing the properties of a device that uses one of these servers, you just have to get to the appropriate field to access the controls to add, edit, or delete these servers.

You can also add these servers if you import them from an inventory file exported from CiscoWorks Common Services Device Credential Repository (DCR) or from another Security Manager server. If you import the server, you bypass the procedure described in this section. For more information about importing devices, see Adding Devices from an Inventory File, page 3-31.

Note Beginning with version 4.18, Cisco Security Manager provides support for ASA 9.10(1) devices that are configured on Umbrella servers. Whenever, the device ID is changed, you must re-discover the device in Cisco Security Manager.

Before You Begin

If you want to populate the Security Manager inventory with your list of AUS and Configuration Engine servers without respect to adding devices, the best approach is to use the New Device wizard and to select **Add New Device** as the add method. This approach is described in this procedure.

You can also add or edit servers by selecting a device in the Device selector and clicking **Tools > Device Properties**. Click **General** in the device properties table of contents. The Server field is in either the Auto Update or Configuration Engine groups. You can add or edit only the type of server identified in the group name.

Tip Security Manager cannot determine the software version running on a Configuration Engine when you add it. However, Security Manager cannot deploy configurations correctly to all versions of Configuration. Ensure that your Configuration Engines are running a supported release (see the release notes for this version of the product to see which Configuration Engine versions are supported at http://www.cisco.com/en/US/products/ps6498/prod_release_notes_list.html).

Related Topics

- Adding Devices from the Network, page 3-12
- Adding Devices by Manual Definition, page 3-26
- Viewing or Changing Device Properties, page 3-41

Step 1 Locate the field that allows you to identify and manage either AUS or Configuration Engine entries in the device inventory:

a. Select **File > New Device** to open the New Device wizard, select **Add New Device** on the Choose Method page, and click **Next**.

b. On the Device Information page, select an ASA device from the Device Type selector, for example, Cisco ASA-5580 Adaptive Security Appliance. The **Server** field in the Auto Update group should include **Add Server** in the drop-down list. It will also include **Edit Server** if there are servers already defined. If these entries have specific server types (for example, Add Auto Update Server or Add Configuration Engine), then you will be limited to adding, editing, or deleting that type of server (in this case, select other types of devices to find the appropriate server type).

Step 2 To add a new AUS or Configuration Engine server, select **Add Server** from the Server drop-down list to open the Server Properties dialog box (see Server Properties Dialog Box, page 3-38).

Step 3 To edit a server, select **Edit Server** from the Server drop-down list to open the Available Servers dialog box (see Available Servers Dialog Box, page 3-39). You can then select the server and click **Edit**, which opens the Server Properties dialog box where you can make your changes.

From the Available Servers dialog box, you can also:

- Click **Create** to add a server.
- Select a server and click **Delete** to remove it from the inventory. You are asked to confirm the deletion. Make sure that the server is not being used by a device in the inventory.

Server Properties Dialog Box

Use the Server Properties dialog box to specify the properties of an Auto Update Server or Configuration Engine.

Depending on how you open this dialog box, the title of the dialog box might specify the type of server (for example, Auto Update Server Properties or Configuration Engine Properties). The dialog boxes are essentially identical.

Tip Security Manager cannot determine the software version running on a Configuration Engine when you add it. However, Security Manager cannot deploy configurations correctly to all versions of Configuration. Ensure that your Configuration Engines are running a supported release (see the release notes for this version of the product to see which Configuration Engine versions are supported at http://www.cisco.com/en/US/products/ps6498/prod_release_notes_list.html).

Navigation Path

To open this dialog box, do one of the following:

- Select **Add Server...** from the **Server** field in the Auto Update Server or Configuration Engine groups on the Device Information page of the New Device wizard when adding a device manually. The selection might also be named Add Auto Update Server or Add Configuration Engine.
- Select **Add Server...** from the **Server** field in the Auto Update Server or Configuration Engine groups on the Device Properties—General page. The selection might also be named Add Auto Update Server or Add Configuration Engine.
- Click **Create**, or select a server and click **Edit**, in the Available Servers dialog box (see Available Servers Dialog Box, page 3-39).

Related Topics

- Available Servers Dialog Box, page 3-39
- Device Information Page—New Device, page 3-27

- Device Information Page – Add Device from Network, page 3-14
- Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36
- Viewing or Changing Device Properties, page 3-41

Field Reference

Table 3-8 Server Properties Dialog Box

Element	Description
Type	The type of server you are defining, either Auto Update Server or Configuration Engine. This field is displayed only if you are adding a server. You cannot change the type of an existing server. For new servers, this field is also not displayed if the title of the dialog box specifies the type of server you are adding.
Server Name	The DNS hostname of the server.
Domain Name	The DNS domain name of the server.
IP Address	The IP address of the server.
Display Name	The name to display in Security Manager for the server.
Username	The username for logging into the server.
Password	The password for accessing the server. In the Confirm field, enter the password again.
Port	The port number that the device managed by the Auto Update Server or Configuration Engine uses to communicate with the server. The port number is typically 443.
URN	This field is displayed only for Auto Update Servers. The uniform resource name for the Auto Update Server. The URN is the name that identifies the resource on the Internet. The URN is part of a URL, for example, /autoupdate/AutoUpdateServlet. The full URL could be: https://: *server ip*:443/autoupdate/AutoUpdateServlet where: • *server ip* is the IP address of the Auto Update Server. • 443 is the port number of the Auto Update Server. • /autoupdate/AutoUpdateServlet is the URN of the Auto Update Server.

Available Servers Dialog Box

Use the Available Servers dialog box to add, edit, or delete an Auto Update Server or Configuration Engine.

Depending on how you open this dialog box, the title of the dialog box might specify the type of servers listed (for example, Available Auto Update Servers or Available Configuration Engines). The dialog boxes are essentially identical.

Each row represents a single server, and shows the display name for the server in Security Manager, its IP address, and DNS hostname and domain name. If the dialog box title does not include the server type, the Type field specifies AUS or CE (Configuration Engine).

- To add a server, click the **Create** button and fill in the Server Properties dialog box (see Server Properties Dialog Box, page 3-38).
- To edit the properties of a server, select it and click the **Edit** button.
- To delete a server, select it and click the **Delete** button. You are asked to confirm the deletion.

Navigation Path

To open this dialog box, do one of the following:

- Select **Edit Server...** from the **Server** field in the Auto Update Server or Configuration Engine groups on the Device Information page of the New Device wizard when adding a device manually. The selection might also be named Edit Auto Update Server or Edit Configuration Engine.
- Select **Edit Server...** from the **Server** field in the Auto Update Server or Configuration Engine groups on the Device Properties—General page. The selection might also be named Edit Auto Update Server or Edit Configuration Engine.

Related Topics

- Device Information Page—New Device, page 3-27
- Device Information Page – Add Device from Network, page 3-14
- Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36
- Viewing or Changing Device Properties, page 3-41

Adding or Changing Interface Modules

Many devices allow you to add or change interface modules. When you make a change to the interface modules hosted in a device, you change the device's inventory.

If you add or change an interface card, you should rediscover the inventory on the device. Rediscovering inventory will replace the Interfaces policy (for routers, the Interfaces > Interfaces policy) and ensure that Security Manager has a correct view of the capabilities of the interfaces available on the device.

Note Inventory rediscovery is especially important for ASA 5580 devices in which you install a 4 GB Ethernet Fiber interface card. For other types of devices, you can usually make manual changes to the Interfaces policy, but rediscovering inventory is the easier and more reliable choice.

Step 1 Right-click the device and select **Discover Policies on Device**.

Step 2 In the Create Discovery Task dialog box, make at least these selections and click **OK** to start rediscovery:

- Discover from **Live Device**.
- Policies to discover: **Inventory**.

Step 3 After discovery completes, edit the Interfaces or Interfaces > Interfaces policy as appropriate and verify that the policy reflects your desired configuration.

Viewing or Changing Device Properties

When you add a device to the inventory, you specify at least some of the device's properties, such as names and credentials. For devices that are in the inventory, you can view and change the device properties.

Related Topics

- Understanding the Device View, page 3-1
- Understanding Device Properties, page 3-6
- Understanding Policies, page 5-1
- Changing Critical Device Properties, page 3-52

Step 1 In Device view, do one of the following in the Device selector to open the Device Properties dialog box:

- Double-click a device.
- Right-click a device and select **Device Properties**.
- Select a device and select **Tools > Device Properties**.

Step 2 In the Device Properties dialog box, click these entries in the table of contents in the left pane to view or change the properties. You must click **Save** before moving from one page to another.

- General—General information about the device, such as the device identity, the operating system running on the device, and transport settings. For information about the fields, see Device Properties: General Page, page 3-41.
- Credentials—The device credentials required to log into the device. For information about the fields, see Device Credentials Page, page 3-45.
- Device Groups—The groups to which the device belongs. For information about the fields, see Device Groups Page, page 3-49.
- Cluster Information—Cluster details for the cluster group, if any. For information about the fields, see Cluster Information Page, page 3-50.
- Policy Object Overrides—The local overrides to policy objects for the device. Policy Object Overrides is a folder that contains the various policy object types that are available for the device. Click a specific policy object type to view the policy objects of that type used by the device and their overrides, if any. For more information about the fields, see Policy Object Override Pages, page 3-52.

Device Properties: General Page

Use the Device Properties General page to add or edit information about the basic properties of the device.

Navigation Path

- From the Device selector, right-click a device and select **Device Properties**, then click **General**.
- From the Device selector, double-click a device, then click **General**.
- Select a device and select **Tools > Device Properties**, then click **General**.

Related Topics

- Understanding Device Properties, page 3-6
- Device Credentials Page, page 3-45
- Device Groups Page, page 3-49
- Policy Object Override Pages, page 3-52

Field Reference

Table 3-9 Device Properties General Page

Element	Description
Identity	
Device Type	The type of device.
IP Type	Whether the IP address for the device is static (defined on the device) or dynamic (supplied by a DHCP server). Depending on the IP type you select, the displayed fields differ. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. The IPv6 address is a 128-bit unique address. For IPv6 address, only Static IP Type is supported. Dynamic IP Type is not supported for IPv6 addresses.
Hostname (Static IP only)	The DNS hostname for the device. This is not necessarily the same name that is configured as the hostname on the device. This property is not updated with the hostname specified in the Hostname device property. It is also not updated with the name defined in the device configuration if you rediscover the device. If you added the device to Security Manager by adding its configuration file, the hostname is initially set to the name specified in the configuration file. If no hostname is specified in the configuration, the name of the file is used as the DNS hostname.
Domain Name (Static IP only)	The DNS domain name for the device.
IP Address (Static IP only)	The management IP address of the device, for example 192.168.3.8. **Note** You must enter either the IP address or the DNS hostname, or both. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. If a device is configured in dual stack, Security Manager would communicate with the device based on the device's IP address added in Security Manager. The IPv6 address is a 128-bit unique address.

Table 3-9 Device Properties General Page (continued)

Element	Description
Display Name	The name to display in the Security Manager Device selector. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; and the following characters: _ -. : and space.
Operating System	
OS Type	The type of operating system. Based on the device type, the OS type is selected automatically. **Note** Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address. This feature is available only for devices where the Operating System type is ASA or FWSM.
Image Name	The name of the image running on the device. The image name is updated whenever you deploy to the device or rediscover its policies.
Running OS Version	The version of the operating system running on the device.
Target OS Version	The OS version on which you want to base the device's configuration. When creating a configuration file using the rules you configure, Security Manager uses commands available in the target OS version. This field is read-only for IPS devices. You cannot change the target OS version to a version that significantly changes the feature set available for the device. For more information, see Changes That Change the Feature Set in Security Manager, page 3-54.
Options	A read-only field whose values are NONE or IPS. The value IPS indicates that the IPS feature is available on the device.
IPS Running OS Version	A read-only field that displays the version of IOS IPS running on the router. This field does not appear if the Options field has the value of NONE.
IPS Target OS Version	A read-only field that displays the target version of IOS IPS running on the router. This field does not appear if the Options field has the value of NONE.
Contexts	Whether the device hosts a single security context (Single) or multiple security contexts (Multi). This field is displayed only if the OS type is an FWSM, ASA, or PIX Firewall 7.0.
Operational Mode	The mode in which the device is operating. This field is displayed only if the OS type is FWSM, ASA, or PIX Firewall 7.0+. The options available are: Transparent or Router. If you choose Multi for Contexts, this mode defaults to Mixed. Mixed applies only to ASA 9.0+ and FWSM 3.1+ devices, and ASA-SMs.

Table 3-9 Device Properties General Page (continued)

Element	Description
Device Communication Settings	
Transport Protocol	The transport protocol that Security Manager should use when accessing the device or deploying configurations to it. If you select **Use Default**, the transport protocol set in the Device Communication page (Tools > Security Manager Administration > Device Communication) is used (see Device Communication Page, page 11-21). You can select a different protocol if the device is not configured to use the default protocol. The available transport protocols differ depending on what the device type supports. For some device types, such as ASA, there is only one option, so the field is grayed out.
CS-MARS Monitoring	
Monitored By	The CS-MARS server that monitors this device, if any. Click **Discover CS-MARS** to have Security Manager determine which CS-MARS server is monitoring the device. If only one CS-MARS server is monitoring it, the field is updated with the server name. If there is more than one, you are prompted to select the CS-MARS server to use. Your selection determines which server is accessed when you try to view CS-MARS collected syslogs or events when viewing firewall access rules or IPS signatures in the policy rule tables for the device. Before you can discover a CS-MARS server for the device, the server must be register with Security Manager on the CS-MARS administration page (Tools > Security Manager Administration > CS-MARS). For more information, see CS-MARS Page, page 11-7.
Auto Update or Configuration Engine This group is named differently depending on the device type: • Auto Update—For PIX Firewall and ASA devices. • Configuration Engine—For Cisco IOS routers. Use these fields to identify the server that manages the device, if any. A server is required for a device with a dynamic IP address.	
Server	The Auto Update Server or Configuration Engine that manages the device. For AUS, this server should match the one defined in the AUS policy (see AUS Page, page 52-1). You can add servers to the list by selecting **Add Servers**, which opens the Server Properties dialog box (see Server Properties Dialog Box, page 3-38. You can also edit the properties of a server by selecting **Edit Server**, which opens the Available Servers dialog box (see Available Servers Dialog Box, page 3-39). For more information on managing this list of servers, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36. For information on how these servers are used during deployment, see Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41.

Table 3-9 Device Properties General Page (continued)

Element	Description
Device Identity	The string value that uniquely identifies the device in Auto Update Server or the Configuration Engine. For AUS, this ID should match the one defined in the AUS policy (see AUS Page, page 52-1).
ASA-CX/FirePOWER Module	
Management IP	The management IP address of the ASA's CX or FirePOWER module; detected during device discovery, or after the module is added to the device. See Detecting ASA CX and FirePOWER Modules, page 72-21 for more information. This field is available only for an ASA CX or FirePOWER module already detected by Security Manager.
Manager Address	The IP address of the Cisco Prime Security Manager (PRSM) or FireSIGHT Management Center used to configure and manage the ASA-CX or FirePOWER module; detected during device discovery, or after the module is added to the device. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information. You can edit this address. However, Security Manager will not perform any validation of the address, and rediscovery or re-detection may alter this address. This field is available only for an ASA CX or FirePOWER module already detected by Security Manager.
Manage in Cisco Security Manager	Whether Security Manager manages the device. Security Manager will not manage configurations nor will it upload or download configurations on this device. You might want to include an unmanaged device in the inventory for these reasons: • If the only function of the device is to serve as a VPN end point. • If the device is a security context that you are using for failover. Because you cannot delete security contexts for managed devices without actually deleting the context from the device itself, you must unmanage the failover contexts.
License Supports Failover (ASA 5505, 5510 only.)	Whether an optional failover license is installed on the device. The option is active for ASA 5505 and 5510 devices only. Security Manager deploys failover policies to the device only if this option is selected. **Tip** If you discover policies from the device, Security Manager determines the license status and sets this option appropriately.

Device Credentials Page

Use the Device Credentials page to add or change the usernames and passwords that are required for device access. For information about device credentials, see Understanding Device Credentials, page 3-4.

The Credentials page is the same whether you are adding a new device (in the New Device wizard), or viewing an existing device's properties.

When adding a new device, you are prompted for credentials only when adding devices manually or from the network.

Tip In the New Device wizard, when you click **Next** or **Finish** when adding a device from the network, Security Manager tests whether it can connect to the device using these credentials. The Device Connectivity Test dialog box stays open while the test is in progress (see Device Connectivity Test Dialog Box, page 9-3). If the test fails, click **Details** to see detailed error information. If you are adding devices that contain modules, for example, a Catalyst switch with an FWSM, you are then prompted for module discovery information.

Navigation Path

- For new devices, to start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.
- For existing devices, to open the device properties, double-click a device in the Device selector, then click **Credentials** on the Device Properties Page.

Related Topics

- Understanding Device Credentials, page 3-4
- Adding Devices from the Network, page 3-12
- Adding Devices by Manual Definition, page 3-26
- Device Communication Page, page 11-21
- Understanding Device Properties, page 3-6
- Viewing or Changing Device Properties, page 3-41
- Managing Device Communication Settings and Certificates, page 9-4
- Discovery Status Dialog Box, page 5-23

Field Reference

Table 3-10 Device Credentials Page

Element	Description
Primary Credentials	
Required for all device types. These credentials are used for SSH and Telnet connections, and for HTTP and HTTPS connections if you select **Use Primary Credentials** in the HTTP group. If you change the password for the specified user, or the enable password, in a device policy, Security Manager uses the old password to log in during deployment. After a successful deployment, the passwords in the device credentials are updated to the newly-deployed passwords. For information on updating the device policies related to these passwords, see the following topics: • ASA/PIX/FWSM devices—Configuring Device Credentials, page 48-17 • IPS devices—Configuring IPS User Accounts, page 36-18 • IOS devices—Defining Accounts and Credential Policies, page 63-14	

Table 3-10 Device Credentials Page (continued)

Element	Description
Username	The user name for logging into the device. The user should have privilege level 15. If the device requires an enable password only to configure it, you can leave the Username and Password fields blank and enter just the Enable Password. **Note** PIX/ASA/FWSM devices require that user names be at least four characters. Passwords can be three to 32 characters; we recommend that passwords be at least eight characters. For ASA devices running the software version 9.6(1) or later, you can enter a password up to 127 characters.
Password	The password for logging into the device (User EXEC mode). In the Confirm field, enter the password again.
Enable Password	The password that activates enable mode (Privileged EXEC mode) on the device if the mode is configured on that device. In the Confirm field, enter the password again.
HTTP Credentials Credentials for making HTTP or HTTPS connections to a device. Some devices support this type of connection, and other devices (such as IPS devices) require it.	
Use Primary Credentials Username Password	Whether Security Manager should use the configured primary credentials for HTTP and HTTPS connections. If the device uses different credentials for HTTP/HTTPS connections, deselect **Use Primary Credentials** and enter the username and password configured for HTTP/HTTPS. Reenter the password in the Confirm field. **Note** PIX/ASA/FWSM devices require that user names be at least four characters. Passwords can be three to 32 characters; we recommend that passwords be at least eight characters. For ASA devices running the software version 9.6(1) or later, you can enter a password up to 127 characters.
HTTP Port	The port to use for HTTP connections. The default is port 80. Change this setting only if the device is configured to accept HTTP connections on a different port.
HTTPS Port	The port to use for HTTPS connections. The default is port 443 (unless a different default is configured in the Security Manager device communication settings). To change the default, first deselect **Use Default**. Change this setting only if the device is configured to accept HTTPS connections on a different port. **Note** If you configure the local HTTP policy to be a shared policy and assign the HTTP policy to multiple devices, the HTTPS port number setting in the shared policy overrides the port number configured in the Device Credentials page for all devices to which the policy is assigned.
IPS RDEP Mode	The connection method to use for contacting IPS devices when making RDEP or SDEE connections (for event monitoring).

Table 3-10 Device Credentials Page (continued)

Element	Description
Certificate Common Name	The name assigned to the certificate. The common name can be the name of a person, system, or other entity that was assigned to the certificate. In the Confirm field, enter the common name again.
Additional Fields and Buttons	
Authentication Certificate Thumbprint (Device properties only.)	The certificate thumbprint for the device that is available in the Security Manager certificate data store. Click **Retrieve From Device** to obtain the current certificate from the device and to replace the one stored in Security Manager. For IPS devices, there are additional options for managing the certificate as described in Managing IPS Certificates, page 44-10.
RX-Boot Mode button	Opens the RX-Boot Mode Credentials Dialog Box, where you can enter the credentials for booting the router from a reduced command-set image (RX-Boot). If these credentials are for a Cisco router that runs from flash memory (where it boots only from the first file in flash), you must run an image other than the one in flash to upgrade the flash image. The RX-Boot credentials are for running this other image.
SNMP button	Opens the SNMP Credentials Dialog Box, where you can specify the SNMP community strings defined on the device.
Test Connectivity button (Device properties and manual device addition only.)	Tests whether Security Manager can connect to the device using the credentials you entered and the configured transport method. For more information about testing device connectivity, see Testing Device Connectivity, page 9-1.

RX-Boot Mode Credentials Dialog Box

Use the RX-Boot Mode Credentials dialog box to add RX-Boot mode credentials, which are used for booting the router from a reduced command-set image (RX-Boot). Enter the RX-Boot Mode username and password; in the Confirm field, enter the password again.

Navigation Path

To open the RX-Boot Mode Credentials dialog box, click **RX-Boot Mode** in the Device Credentials Page in either the New Device wizard (when adding a device manually or from the network), or the Device Properties page.

SNMP Credentials Dialog Box

Use the SNMP Credentials dialog box to add SNMP credentials.

Navigation Path

To open the SNMP Credentials dialog box, click **SNMP** in the Device Credentials Page, page 3-45 in either the New Device wizard (when adding a device manually or from the network), or the Device Properties page.

Field Reference

Table 3-11 SNMP Credentials Dialog Box

Element	Description
SNMP V2C	
These are the credentials for devices running SNMP version 2.	
RO Community String	The read-only community string. In the Confirm field, enter the community string again.
RW Community String	The read-write community string. In the Confirm field, enter the community string again.
SNMP V3	
These are the credentials for devices running SNMP version 3.	
Username	The SNMP version 3 authentication user name.
Password	The SNMP version 3 authentication user password. In the Confirm field, enter the password again.
Auth Algorithm	The authorization algorithm for encrypting the password. You can choose MD5 or SHA-1.
Privacy Password	The SNMP version 3 encryption user password. In the Confirm field, enter the password again.
Privacy Algorithm	Specify the encryption level by choosing an encryption algorithm and version: • DES – Apply the Data Encryption Standard cipher algorithm, using 56-bit keys.. • 3DES – Use Triple DES; the Data Encryption Standard cipher algorithm is applied three times to each packet. • AES128 – Use the Advanced Encryption Standard with 128-bit keys. • AES192 – Use the Advanced Encryption Standard with 192-bit keys. • AES256 – Use the Advanced Encryption Standard with 256-bit keys.
Engine ID	Enter the hexadecimal identifier for the SNMP v3 authorization agent in the device.

Device Groups Page

Use the Device Groups page to assign the device to device groups. You can also edit or delete device groups from this page.

Navigation Path

- For new devices, to start the New Device wizard, from Device view, select **File > New Device**, or click the **Add** button in the device selector.
- For existing devices, to open the device properties, double-click a device in the Device selector, then click **Device Groups** on the Device Properties Page.

Related Topics

- Understanding Device Grouping, page 3-60
- Adding Devices to the Device Inventory, page 3-7
- Understanding Device Properties, page 3-6
- Discovery Status Dialog Box, page 5-23

Field Reference

Table 3-12 Device Grouping Page

Element	Description
Group Types, such as Department and Location	The group types defined in Security Manager, for example, Department or Location. Each field contains a list of the device groups defined within that group type. Select the device groups to which the device should belong. If you want to create a new device group, or group type, select **Edit Groups** from the drop-down list for any of the existing group types. This opens the Edit Device Groups page, where you can create new groups and group types or delete them (see Edit Device Groups Dialog Box, page 3-61).
Set values as default	Whether to set the selected groups as the default groups. If you select this option, other devices you add are automatically added to these groups.

Cluster Information Page

Use the Device Properties Cluster Information page to view details for a cluster.

Navigation Path

- From the Device selector, right-click a device and select **Device Properties**, then click **Cluster Information**.
- From the Device selector, double-click a device, then click **Cluster Information**.
- Select a device and select **Tools > Device Properties**, then click **Cluster Information**.

Related Topics

- Working with Device Clusters, page 3-9
- Understanding Device Properties, page 3-6
- Device Credentials Page, page 3-45
- Device Groups Page, page 3-49
- Policy Object Override Pages, page 3-52

Field Reference

Table 3-13 Device Properties Cluster Information Page

Element	Description
Cluster Details	
Device Type	The type of device.
Cluster Group Name	The name assigned to the cluster.
Cluster Master	The cluster member name of the device that is serving as the control unit. **Note** Changes to the control unit are not automatically reflected in Security Manager.
Retrieve From Device	Use **Retrieve From Device** to update the control unit information.
Interface Mode	Whether the interfaces are configured for Layer 2 load balancing (Spanned EtherChannel) or Layer 3 load balancing (Individual).
Management IP Pool Range	Enter the IP address pool used for cluster management. You can provide this value for the device in user context. This field is mandatory, if Eventviewer is being used to monitor syslogs for a Multi-context ASA cluster. If you leave this field blank, or enter an incorrect IP address pool, Eventviewer cannot categorize the syslogs for a specific context and drops the syslog events. **Note** Ensure that you enter valid IP addresses. Cisco Security Manager will not be validating the entered IP address pool.
Last Update in CSM	The date and time that cluster information was last updated for this cluster by Security Manager.
Cluster VPN Mode	Beginning from Cisco Security Manager 4.16, after discovering the Cluster device, cluster VPN mode will be displayed. This value will be Centralized or Distributed. **Note** This value is also displayed in the pop up window that appears when you hover the mouse pointer over the device in the device selector view.
Cluster VPN Backup	Beginning from Cisco Security Manager 4.16, the Cluster VPN Backup is displayed. One of the following values will be displayed for distributed mode — • Flat —When cluster VPN backup is on any other member • Remote Chassis — When cluster VPN backup is on a different chassis Cluster VPN Backup information is not shown for centralized VPN mode. The value for this field in a centralized VPN mode is N/A. **Note** This value is also displayed in the pop up window that appears when you hover the mouse pointer over the device in the device selector view.

Cluster Node Details

The Cluster Node Details table lists details for each device in the cluster.

Table 3-13 Device Properties Cluster Information Page (continued)

Element	Description
Cluster ID	The cluster ID of the cluster node.
Node Name	The member name of the cluster node.
Serial Number	The serial number of the cluster node.
CCL IP	The cluster control link IP address for the cluster node.
CCL MAC	The cluster control link MAC address for the cluster node.
Site ID	The site that the current cluster member belongs to. Configuring a site ID prevents MAC address flapping.

Policy Object Override Pages

You can override the global settings for many types of policy objects from the Device Properties window of a selected device. This enables you to customize the definition of an object on that device. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18.

The Policy Object Overrides folder in the table of contents includes all of the types of objects for which you can create overrides for the particular type of device. When you select an object type, the existing policy objects that are configured to allow device overrides appear in the table in the right pane, if any. If an object has an override already defined for the device, the Value Overridden? column contains a check mark.

You can create and manage overrides for these objects. Select an object and you can do the following:

- To create an override, click the Create Override button. This opens the edit dialog box for that type of object. Click the Help button for object-specific information.
- To edit an existing override, click the Edit Override button.
- To remove an override, click the Delete Override button.

Navigation Path

Double-click a device in the Device selector, then click the desired policy object type in the **Policy Object Overrides** folder in the table of contents in the left pane.

Related Topics

- Policy Object Overrides Window, page 6-20
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19
- Deleting Device-Level Object Overrides, page 6-21
- Filtering Tables, page 1-48

Changing Critical Device Properties

You must use caution when changing the image version of a device, the device type, or the security context or operational mode of FWSM and ASA devices that are managed by Security Manager. In certain cases, these changes enable a different set of features for the device. As a result, some of the policies that you configured for the device in Security Manager might no longer apply.

The key device changes, their effect on the policies available in Security Manager, and the procedure you should follow to implement these device changes, are described in the following sections:

- Image Version Changes That Do Not Change the Feature Set in Security Manager, page 3-53
- Changes That Change the Feature Set in Security Manager, page 3-54

Image Version Changes That Do Not Change the Feature Set in Security Manager

The following image version changes *do not* affect the types of policies available for that device in Security Manager:

- Upgrading from one IOS individual release number to another individual release number within the same Cisco IOS release; for example, upgrading from IOS 12.3(10) to 12.3(13).
- Upgrading from any IOS 12.1 image to any 12.2 image.
- Upgrading from any IOS 12.2 image to any 12.3 image.
- Upgrading from any IOS 15.0 image to any 15.1 image.
- Upgrading from any IOS 15.2 image to any 15.3 image.
- Upgrading from any PIX 6.x image to another PIX 6.x image.
- Upgrading from any PIX 7.x image to another PIX 7.x image, retaining the same security context and mode configuration.
- Upgrading from any ASA 7.x image to another ASA 7.x image, retaining the same security context and mode configuration.
- Upgrading from any ASA 8.0(x)-8.2(x) image to another ASA 8.0(x)-8.2(x) image, retaining the same security context and mode configuration.
- Upgrading from any FWSM 2.x image to another 2.x FWSM image, retaining the same security context and mode configuration.
- Upgrading from any FWSM 3.x image to another 3.x FWSM image, retaining the same security context and mode configuration.
- Upgrading a Catalyst 6500/7600 chassis from any IOS 12.x image to another IOS 12.x image.

Note This list applies only to images that are supported by Security Manager. For a list of supported images, see *Supported Devices and Software Versions for Cisco Security Manager* for this version of the product at http://www.cisco.com/en/US/products/ps6498/products_device_support_tables_list.html.

For these cases, use the following procedure to change the image version.

Related Topics

- Understanding the Device View, page 3-1
- Understanding Device Properties, page 3-6
- Understanding Policies, page 5-1
- Changes That Change the Feature Set in Security Manager, page 3-54

Step 1 Upgrade the image version on the device.

Step 2 In Device view, do one of the following in the Device selector to open the Device Properties dialog box:

- Double-click a device.

- Right-click a device and select **Device Properties**.
- Select a device and select **Tools > Device Properties**.

Step 3 In the Device Properties dialog box, change the **Target OS Version** property on the General page to the updated version number and click **Save**.

Changes That Change the Feature Set in Security Manager

These are the main types of device changes that affect the policy feature set available for a device:

- Image version changes—The following image version changes affect the types of policies available for that device in Security Manager:
  - Upgrading to ASA 8.4(x) or higher from an ASA 8.3(x) or lower release.
  - Upgrading to ASA 8.3(x) or higher from an ASA 8.2(x) or lower release.
  - Changes in the major version number for ASA, PIX, FWSM, and IPS devices. For example, upgrading an ASA from 8.x to 9.x, or downgrading an IPS device from 7.x to 6.x.
  - Upgrading from an IOS 12.1 or 12.2 image to an IOS 12.3 or 12.4 image.
  - Downgrading from an IOS 12.3 or 12.4 image to an IOS 12.1 or 12.2 image.
  - Upgrading to IOS 15.2 or higher from an IOS 12.3 or lower release.

  If you make these changes, and you do not have any policies defined that are affected by the change, you might be able to change the target OS version of the device. Security Manager prevents you from changing the target OS version of a managed device to a version that changes the types of policies that are available for that device, and informs you when it cannot make the change (identifying the problem policies). Therefore, you must first delete the device from Security Manager, perform the image change, then add the device back.

  Certain types of policies, such as access rules, are not affected by changes in image version or changes in platform type.

  Changes to NAT policies that were introduced in the 8.3 and 9.0.1 ASA releases require that the NAT policies are rediscovered in Security Manager. This can be accomplished by deleting the device and then adding it back in to Security Manager, as described below, or you can rediscover just the NAT policies using the Discover Policies on Device feature. For more information on the Discover Policies on Device feature, see Discovering Policies on Devices Already in Security Manager, page 5-15.

  Note If an ASA device was upgraded or downgraded from the current version to a higher or lower version outside of Security Manager, you should delete the device and then add it back in to Security Manager.

- Security context and operational mode changes—Changes that you make to the security context and operational mode settings on an FWSM or ASA device enable a different set of features on that device. These changes occur if you change the device from:
  - Single context to multiple context (or vice-versa).
  - Routed mode to transparent mode (or vice-versa).

Security Manager prevents you from changing the security context or operational mode settings of a managed device. Therefore, you must first delete the device from Security Manager, change the context or mode, then add the device back.

Certain policy types (for example, Banner, Clock, Console Timeout, and HTTP) are not affected by changes in operational mode. Other policy types (for example, ICMP, SSH, and TFTP, in addition to Banner and Clock) are not affected by changes in security context settings.

- Replacing device hardware—In some cases, you might replace a particular device but retain the original contact information (such as the IP address), for example:
  - Replacing a PIX firewall with a Cisco IOS router.
  - Replacing a PIX firewall with an ASA device.
  - Replacing a router with a firewall device.
  - Replacing a router with a new router of a different model.

  In all of these cases, the new device changes the types of policies available for that device in Security Manager. Security Manager prevents you from modifying the hardware model of an existing device. Therefore, you must first delete the device from Security Manager, change the physical device, then add the device back.

  Certain policy types (for example, access rules) are not affected by changes in device type.

We recommend that you share the policies configured on your device that will not be affected by the change before you remove it from Security Manager. This provides a useful method for reassigning the policies to the device (with any inheritance and policy object references intact) after you add it back to Security Manager. The following procedure describes how to do this.

Related Topics

- Understanding the Device View, page 3-1
- Understanding Device Properties, page 3-6
- Understanding Policies, page 5-1
- Image Version Changes That Do Not Change the Feature Set in Security Manager, page 3-53

Step 1 Submit and deploy all the changes you configured for the device in Security Manager. This ensures that the desired configuration is on the device before the image upgrade.

Step 2 Share the local policies defined on the device:

a. Right-click the device in the Device selector, then select **Share Device Policies**. By default, all policies configured on the device (local and shared) are selected for sharing in the Share Policies wizard.

b. Deselect the check box next to each existing shared policy, as indicated by the hand in the policy icon. You should do this because there is no need to create a copy of the shared policies that already exist; you will reassign the existing shared policies after the image version upgrade.

c. Enter a name for the shared policies. We recommend using the device name as a convenient means of identification. For example, if the device name is MyRouter, each shared policy is given the name MyRouter. Make a note of all the policies you are creating for this purpose.

d. Click **Finish**. The selected local policies become shared policies.

Step 3 Delete the device from Security Manager.

Step 4 Make the desired change to the device, for example, upgrade the image version, change the operational mode, or replace the device.

Step 5 Add the device back to Security Manager and perform policy discovery.

Step 6 Reassign the policies to the device:

a. Right-click the first policy type displayed in the Device Policies selector, then select **Assign Shared Policy**.

b. In the Assign Shared Policy dialog box, do one of the following:

- If a local policy was previously defined on the device, select the shared policy you created for this procedure and click **OK**.
- If a shared policy of this type was previously assigned to the device, select it and click **OK**.

c. (Local policies only) Right-click the policy type again in the Device Policies selector, then select **Unshare Policy**.

d. Repeat the process for each policy type that is relevant to the device's configuration. If a shared policy is not available, this indicates that this is a policy type that was not available for the previous image version.

Step 7 (Optional) Delete the shared policies created for this procedure from Policy view:

a. Select **View > Policy View** or click the **Policy View** icon on the toolbar.

b. Select one of the policies you want to delete and click the **Assignments** tab in the work area to verify that the policy is not assigned to any devices.

c. Click the **Delete Policy** button beneath the Shared Policy selector to delete the policy.

d. Repeat the process for each policy type that you want to delete.

Showing Device Containment

You can display the service modules, security contexts, and virtual sensors that are contained in devices that include them. Based on the type of device, you can view these contained elements:

- Catalyst 6500 devices—The IDSM and FWSM service modules, security contexts, and virtual sensors.
- For FWSM, PIX Firewall 7.0, and ASA devices—The security contexts defined on the device. For information about security contexts, see Chapter 59, "Configuring Security Contexts on Firewall Devices".
- IPS devices—The virtual sensors defined on the device.

To view contained items, in Device view, select one of these types of devices and then select **Tools > Show Containment**, or right-click the device and select **Show Containment**. The Composite View dialog box opens and displays elements contained in the selected device, if any.

Cloning a Device

A cloned (duplicate) device shares the configurations and properties of the source device. Cloning a device saves you time because you do not need to re-create configuration and properties on the new device.

The cloned device shares the device operating system version, credentials and grouping attributes with the source device, but it has its own unique identity, such as display name, IP address, hostname, and domain name. You can clone only one device at a time.

Note You cannot clone a Catalyst switch or a Catalyst 6500/7600 device.

Related Topics

- Understanding the Device View, page 3-1
- Copying Policies Between Devices, page 5-33

Step 1 Do one of the following:

- (Device view) Select the device and select **File > Clone Device**, or right-click the device in the Device selector and select **Clone Device**.
- (Map view) Right click a device and select **Clone Device**.

The Create a Clone of Device dialog box appears.

Step 2 Enter the IP address and names for the clone in the appropriate fields. Following are the available attributes:

- **IP Type**—Whether the device uses a static or dynamic (DHCP-provided) IP address. You cannot change the IP type when cloning a device.
- **Hostname**—(Static IP only.) The DNS hostname for the cloned device.
- **Domain Name**—(Static IP only.) The DNS domain name for the cloned device. If you do not provide the domain name, Security Manager uses the default domain name configured on the server.
- **IP Address**—The management IP address of the cloned device, for example, 10.10.100.1. If you do not know the IP address, enter the DNS hostname in the Hostname field. You must enter either the IP address or the hostname for devices with static IP addresses.

Note Beginning with version 4.12, Security Manager server to device communication for ASA devices is supported over either IPv6 address or over IPv4 address.

- **Display Name**—The name that appears in Security Manager device lists. The maximum length is 70 characters. Valid characters are: 0-9; uppercase A-Z; lowercase a-z; and the following characters: _ -. : and space.
- **Device Identity**—(Dynamic IP only.) The string value that uniquely identifies the device in Auto Update Server or Configuration Engine. This field appears only if the device is configured to use one of these servers.
- **Clone VPN Assignments**—Whether to copy the VPN assignments defined for the device. This field is displayed only if the device supports VPN assignments.

  You can clone the VPN assignments of a device that is a spoke in a hub-and-spoke configuration, or a device that participates in a full mesh topology. If you clone a spoke device, the new device is added to the VPN as a new spoke with the same policies. If you clone a device in a full mesh VPN, the new device is added to the full mesh VPN with the same policies. You cannot clone a device in a point-to-point VPN topology.

Step 3 Click **OK**. A clone of the source device with its unique display name is created in the Device selector.

Deleting Devices from the Security Manager Inventory

If you do not want to continue managing a device in Security Manager, you can delete it from the inventory. Deleting a device from Security Manager does not change any configuration settings on the device.

Tip If someone is configuring policies on the device, locks will prevent you from deleting the device.

There are special considerations when deleting certain types of devices:

- If the device participates in a VPN, deleting the device removes it from the VPN. However, if removing the device invalidates the VPN topology, the entire VPN topology is also deleted when you delete the device. You are warned of this and given the opportunity to cancel the device deletion.
- For ASA, PIX, and FWSM devices running in multiple context mode, or for IPS devices that contain virtual sensors, deleting the device also deletes all of its security contexts or virtual sensors. You cannot delete an individual security context or virtual sensor using this procedure: instead, you must modify the appropriate policies on the hosting device to remove them.
- If you delete a device that contains managed service modules, the contained devices are also deleted. For example, if you added a Catalyst switch and its contained FWSM, if you delete the switch, the FWSM is also deleted. You are warned if contained devices will be deleted.

Tip Device deletion requires the removal of a lot of information from the database. If you delete a lot of devices at one time, it can take a while for the operation to complete. If you have a lot of devices to delete, consider deleting them in smaller groups.

Step 1 In Device view, do one of the following:

- Select the devices you want to delete, or a device group if you want to delete all devices within the group, right-click and select **Delete Devices**. You can also click the **Delete Device** button (the trash can icon) above the device selector.
- Select **File > Delete Device**, then select the devices to delete in the Device Selector dialog box and click **>>** to move them to the selected devices list (which is pre-filled with any devices that were selected in the device tree). You can select a device group to delete all of its member devices. Click **OK** when finished.

Tip When you select a device group, you are deleting only the devices in the group, you are not deleting the group itself. For information on deleting device groups, see Deleting Device Groups or Group Types, page 3-63.

Step 2 You are asked to confirm that you want to delete the devices.

Security Manager then validates whether the device can be deleted. If problems or potential problems are identified, they are listed in the Device Delete Validation Dialog Box, page 3-59. This dialog box shows errors (indicating devices that cannot be deleted) as well as warnings and informational messages.

You can elect to confirm the deletion of devices that have warnings or informational messages if you accept the consequences described in the message. The dialog box has an **OK** button if you can continue deleting all selected devices, or a **Continue** button if there are any error messages. If you click **Continue**, you are deleting only those devices without error conditions. You are asked to confirm.

Device Delete Validation Dialog Box

Use the Device Delete Validation dialog box to view error, warning, and informational messages during device deletion. For detailed information on deleting devices, see Deleting Devices from the Security Manager Inventory, page 3-58.

Each row represents a device for which a validation issue arises when trying to delete it. Displayed are the message severity icon, device display name, and the result of validation, which indicates the reason why you cannot delete the device, or warnings or information about the perhaps unexpected consequences of deleting the device. If there are no messages for a device, it is not listed.

Double-click a row or select it and click the **Details** button to read longer messages. The information is displayed in the Device Delete Validation Details dialog box in a more readable format.

The message severity can be one of the following.

- Error—A problem was detected that will prevent you from deleting the device. For example, another user has a lock on a device.
- Warning—Proceed with caution. For example, deleting the device will invalidate a VPN topology, and if you continue, the VPN topology will also be deleted.
- Information—A minor problem exists. For example, deleting the device will delete it from a VPN.

To proceed with the device deletion, click the OK or Continue button, which is actually the same button:

- If the text says **OK**, then when you click it, all devices you selected for deletion are deleted.
- If the text says **Continue**, then there are errors for some of the devices you selected. If you click Continue, you will delete only those devices that do not have errors.

If all selected devices have errors, the button is greyed out and you must click Cancel. Resolve any errors before attempting to delete the devices.

Navigation Path

This dialog box appears only if you try to delete devices and Security Manager determines that there are problems with the deletion.

Working with Device Groups

You can create device groups to help you organize your devices for more effective device management. The following topics explain device groups and how to use them:

- Understanding Device Grouping, page 3-60
- Creating Device Group Types, page 3-62
- Creating Device Groups, page 3-62
- Deleting Device Groups or Group Types, page 3-63
- Adding Devices to or Removing Them From Device Groups, page 3-63

Understanding Device Grouping

Device groups are simple, arbitrary, organizational collections of devices that you create for more effective network visualization. They are not policy-sharing entities. They are distinct from the various policy object groups (for example AAA server group objects and user group objects). For information on policy objects, see Chapter 6, "Managing Policy Objects".

Tip If you have a large number of devices, grouping them can make it easier to select a subset of devices when you deploy changes to them. For example, if there is a set of devices that you know you will want to deploy changes to simultaneously, if you put them in a single device group, you just have to select the group in the deployment job. For more information about policy deployment, see Chapter 8, "Managing Deployment".

Device grouping enables you to view a subset of devices in the inventory. The device group hierarchy has two types of folders:

- Device group types—Group types are the highest level in the hierarchy. A group type can contain specific device groups, but it cannot contain devices, except for the All group type, which includes all devices in the inventory. Security Manager comes with the group types Department and Location predefined, but you do not need to use them, and you can delete them. You can create a maximum of 10 group types.
- Device groups—Device groups are subfolders within a group type folder. You can create multiple levels of nested device groups. You can place devices within device groups. However, a device can be in only one group within a group type. For example, in Figure 3-3 under the group type, Location, you can assign routerx to San Jose, but you cannot assign routerx to San Jose and California.

Figure 3-3 shows an example of nested device groups with devices in some of the groups. Notice that an individual device can reside in multiple groups. In this example, routerx is in the Finance group (under the Department group type), and also in the Location > United States > California > San Jose nested group. If you select routerx in any of these places, you are configuring a single device (the configurations are not tied to the grouping).

Figure 3-3 Device Groups

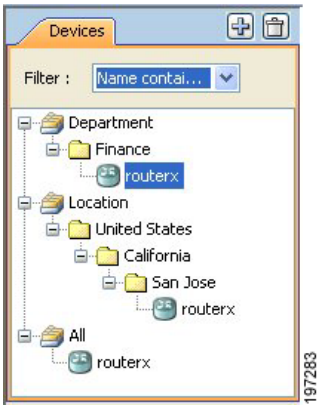

Security Manager lets you create or delete group and group types, and put devices in groups, in many locations in the interface:

- When adding devices to the inventory—The New Device wizard includes a Device Grouping page, where you can create device group types and select a group for the newly-added device. You can also select a default group to which all new devices are added.
- When viewing the device inventory in Device view—The File > Edit Device Groups command opens a dialog box where you can create or delete groups and group types. If you select a group or group type in the Device selector, the File menu and the right-click shortcut menu includes commands for adding groups or adding devices to groups.

  To add devices to a group, or remove them from a group, select the group and select **File > Add Devices to Group**.
- When viewing the properties for a device—The Device Grouping page allows you to select the groups to which the device belongs, and to set defaults for devices added to the inventory. This is the only place where you can remove a device from a device group. Double-click a device in the Device selector to open the device properties.
- When using the administration pages—Select **Tools > Security Manager Administration > Device Groups** to open the administration page for device groups, where you can create or delete groups and group types, but you cannot add devices to groups here.

Related Topics

- Creating Device Group Types, page 3-62
- Creating Device Groups, page 3-62
- Deleting Device Groups or Group Types, page 3-63
- Adding Devices to or Removing Them From Device Groups, page 3-63

Edit Device Groups Dialog Box

Use the Edit Device Groups dialog box to manage the device groups and group types defined in the device inventory.

Navigation Path

Do one of the following:

- Right-click a device group type or a device group in the Device selector and select **Edit Device Groups**.
- Select **File > Edit Device Groups**.
- From the Device Grouping page in the New Device wizard or for existing devices, the device properties, select **Edit Groups** from a group type list. See Device Groups Page.

Related Topics

- Understanding Device Grouping, page 3-60
- Working with Device Groups, page 3-59

Field Reference

Table 3-14 Edit Device Groups Dialog Box

Element	Description
Groups	Displays the device groups and group types. To rename a group or type, select it and then click it again to make the text editable. Type in the new name and press Enter.
Add Type button	Click this button to create a new group type. The type is added with a default name. Overtype the name and press Enter. You can have a maximum of 10 group types.
Add Group to Type button	Click this button to add a device group to the selected device group or group type.
Delete button (trash can)	Click this button to delete the selected device group or group type and all device groups that it contains. Deleting a device group or group type does not delete any devices it contains.

Creating Device Group Types

This procedure describes the most direct method to create device group types. For information on other methods of adding group types, see Understanding Device Grouping, page 3-60.

Device group types are the top-level categories in your device group hierarchy. If you want add a device group, see Creating Device Groups, page 3-62

Related Topics

- Understanding Device Grouping, page 3-60
- Deleting Device Groups or Group Types, page 3-63
- Adding Devices to or Removing Them From Device Groups, page 3-63

Step 1 Select **File > Edit Device Groups**.

The Edit Device Groups page opens (see Edit Device Groups Dialog Box, page 3-61).

Step 2 Click **Add Type**. A new device group type entry is added to the selector.

Step 3 Enter a name for the group type and press **Enter**.

Step 4 Click **OK** to close the Edit Device Groups page.

Creating Device Groups

This procedure describes the most direct method to create device groups. For information on other methods of adding groups, see Understanding Device Grouping, page 3-60.

Device groups are the lower-level categories in your device group hierarchy, and are added either within a device group type (top-level) or within another device group. If you would rather add a device type group, see Creating Device Group Types, page 3-62.

Related Topics

- Understanding Device Grouping, page 3-60
- Adding Devices to or Removing Them From Device Groups, page 3-63
- Deleting Device Groups or Group Types, page 3-63

Step 1 Select a device group or group type in the Device selector and select **File > New Device Group**, or right-click and select **New Device Group**.

The Add Group dialog box appears.

Step 2 Enter a name for the device group and click **OK**. The new device group is added to the Device selector.

Deleting Device Groups or Group Types

If you no longer need a device group or group type, you can delete it. The only group type that you cannot delete is the All group.

When you delete a group or group type, you delete any groups that are in it. However, you are not deleting any devices. The devices that are in the group remain in the inventory and can be found in other groups to which they belong (you can find all devices in the All group).

There are many ways to delete device groups and group types. This procedure explains the most direct way. For information on other methods of deleting them, see Understanding Device Grouping, page 3-60.

Step 1 In Device view, select **File > Edit Device Groups**. The Edit Device Groups page opens (see Edit Device Groups Dialog Box, page 3-61).

Step 2 Select the group type or group you want to delete and click the **Delete** button. You are asked to confirm the deletion.

Adding Devices to or Removing Them From Device Groups

You must create a device group before you add devices to it. To create groups, see Creating Device Groups, page 3-62.

Related Topics

- Understanding Device Grouping, page 3-60
- Filtering Items in Selectors, page 1-45

Step 1 Select the device group in the Device selector, right-click and select **Add Devices to Group**. The Add Devices to Group dialog box appears.

Step 2 To add devices to the group, select the devices in the Available Devices selector and click **>>** to move them to the Selected Devices list.

To remove devices, select them in the Selected Devices list and click **<<**.

Step 3 Click **OK**. The device group membership is adjusted to include the devices that were in the Selected Devices list.

Working with Device Status View

You can use the Device Status View to quickly see the status of the devices in the Security Manager inventory. The Device Status View window aggregates information from several applications and tools within Cisco Security Manager. You can use the Device Status View to quickly see the status of all your devices or specific groups of devices and can easily navigate to the areas in Security Manager you need to act on that information.

Caution In some cases, for a particular device, Health and Performance Monitor displays a "critical" device status while Configuration Manager displays a "normal" device status. Restarting the services or the server does not resolve this discrepancy. For this reason, you should monitor device status in HPM in addition to Configuration Manager.

Navigation Path

- Select **View > Device Status View**. The Device Status View window opens showing information for all devices.
- Select a device group in the Device selector. The Device Status View window opens showing information for the devices that are part of that device group or a subgroup.

Figure 3-4 Device Status View

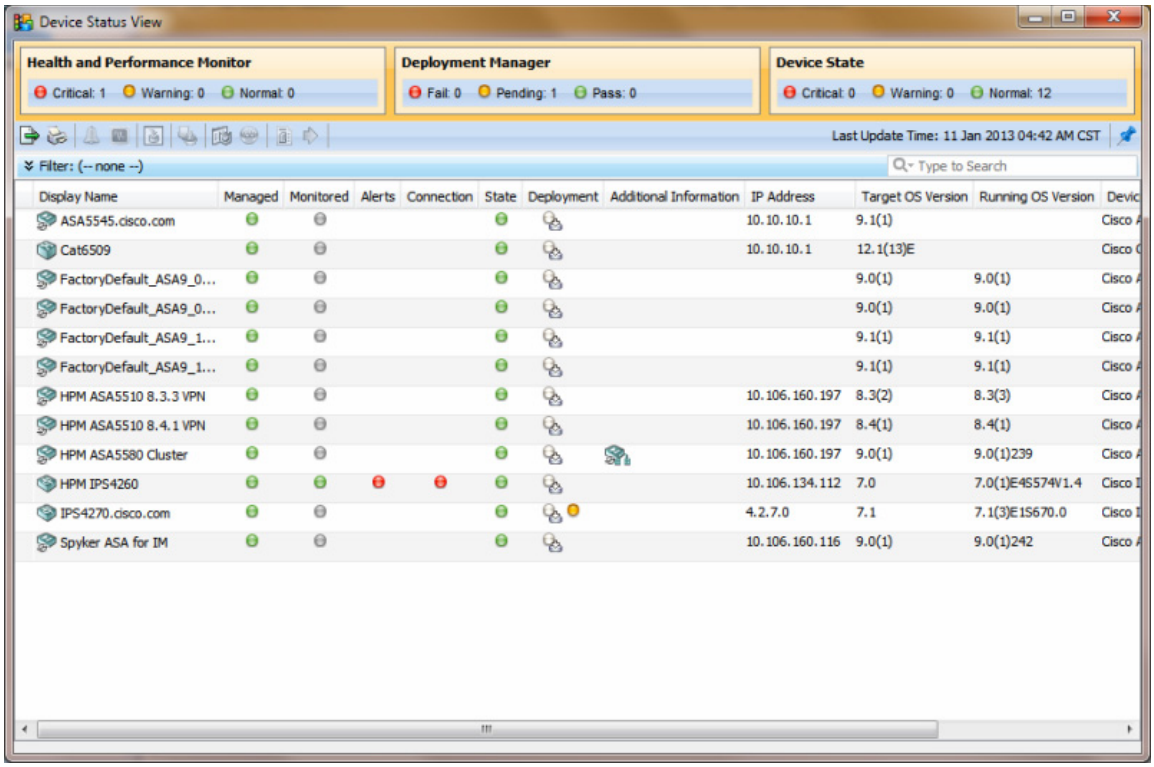

Field Reference

Table 3-15 Device Status View

Element	Description
Device Status Summary Boxes	
The Device Status Summary boxes provide a high-level view of the overall status of the devices in the Device Status View. The counts shown in the summary boxes reflect the status for the devices in the currently selected device group. If you select View > Device Status View or select the All devices group, then the summary boxes reflect the counts for all devices. **Note** Filtering the device list in the Device Status View window will not affect the counts in the Device Status Summary boxes.	
Health and Performance Monitor summary box	Shows the device counts for the Critical (red), Warning (yellow), and Normal (green) alert statuses.
Deployment Manager summary box	Shows the device counts for the Fail (red), Pending (yellow), and Pass green) deployment statuses.
Device State summary box	Shows the device counts for the Critical (red), Warning (yellow), and Normal (green) device states.

Table 3-15 Device Status View (continued)

Element	Description
Device Status View Toolbar	
The Device Status View toolbar provides the following buttons:	
Note These options are all also available from the right-click menu for a device.	
	Allows you to export the device status information to a PDF file.
	Allows you to print the device status information.
	Shows alert status information for the selected device in the Health & Performance Monitor application. For more information, see Chapter 71, "Health and Performance Monitoring".
	Shows monitoring information for the selected device in the Health & Performance Monitor application. For more information, see Chapter 71, "Health and Performance Monitoring".
	Opens the Deployment Manager. For more information, see Chapter 8, "Managing Deployment".
	Opens the Image Manager application for the selected device. For more information, see Chapter 73, "Using Image Manager".
	Opens the device manager for the selected device. For more information, see Starting Device Managers, page 72-14.
	Launches the Cisco Prime Security Manager (PRSM) application for the selected device. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.
	Opens the Device Properties dialog box for the selected device. For more information, see Viewing or Changing Device Properties, page 3-41.
	Allows you to navigate to the selected device from the Device Status View window. For more information, see Understanding the Device View, page 3-1.
	Opens online help for the current page. For more information, see Accessing Online Help, page 1-52.
	Undocks the Device Status View window, which enables you to use other product features while keeping the window open.
	Docks the Device Status View window. **Note** If the selection has changed in the Device Selector, the Work area will reflect the current selection when the Device Status View window is docked.
Table Filter	
You can filter the list of devices displayed in the Device Status View table to help you find items meeting specific criteria. For more information, see Filtering Tables, page 1-48.	

Table 3-15 Device Status View (continued)

Element	Description
Device Status Table	
Display Name	The display name for the device. This is the name used for display in the Security Manager Device selector and is not necessarily the same as the host name for the device.
Managed	Whether Security Manager manages the device.
Monitored	Whether the device is monitored by the Health and Performance Monitor.
Alerts	Indicates current alert level for the device; can be Normal (green), Warning (yellow), or Critical (red). You can hover over the alert indicator to view more details.
Connection	Indicates HPM's ability to connect to/poll the device: Connected, Authentication Error, Certificate Mismatch Error, Connection error, Timeout during Read operation, or Service unavailable. You can hover over the alert indicator to view more details. **Note** If the device is not selected as a Normal or Priority Monitored Device in HPM (Tools > Device Selector), this status will not apply. Changes to Monitored Device selection may take several minutes to become effective and be reflected on screen.
State	Indicates the current state of the device. You can hover over the alert indicator to view more details. For ASA devices that are being monitored by the Health and Performance Monitor, the State column will also alert when possible out of band changes have been detected. Any out of band changes that occurred prior to monitoring the device in Health and Performance Monitor will not be reflected in the State column. For more information about out of band changes, see Understanding How Out-of-Band Changes are Handled, page 8-12 and Detecting and Analyzing Out of Band Changes, page 8-45.
Deployment	Indicates the deployment method and the current deployment status for the device. Deployment status can be Fail (red), Pending (yellow), and Pass (green). You can hover over the alert indicator to view more details.
Additional Information	Shows additional information for the device, such as whether the device is in cluster mode. You can hover over the alert indicator to view more details.
IP Address	The management IP address of the device, for example 192.168.3.8.
Hostname.Domain	The DNS hostname and domain name for the device.
Target OS Version	The OS version on which you the device's configuration is based.
Running OS Version	The version of the operating system running on the device.
Device Type	The type of device.

Related Topics

- Chapter 71, "Health and Performance Monitoring"
- Chapter 73, "Using Image Manager"
- Chapter 8, "Managing Deployment"

Managing Activities

Whether you are using Workflow or non-Workflow mode, all policy configuration is done within an activity, which is also called a configuration session in non-Workflow mode. In Workflow mode, you must explicitly create and manage activities, whereas in non-Workflow mode much of the activity creation and management is done automatically for you. However, in non-Workflow mode, you are in fact working within an activity whenever you modify policies, and so you should understand the basic activity concepts.

In non-Workflow mode with Ticket Management enabled, an activity is created automatically and transparently whenever you open a ticket. You must actively open and manage tickets in this mode.

The following topics provide information about activities:

- Understanding Activities, page 4-1
- Working with Activities/Tickets, page 4-7

Understanding Activities

An activity is a temporary context within which you define policies and assign them to devices. You do not need to create an activity to import, create, or delete devices (unless you perform policy discovery as part of the action), or to perform various system management tasks.

The requirements for creating or opening activities differ depending on your Workflow mode:

- Non-Workflow mode with Ticket Management—An activity is created automatically and transparently whenever you open a ticket. If you do not explicitly open a ticket, you are prompted to create a new ticket or open an existing one whenever you perform an action that requires an activity. You must actively open and manage tickets in this mode.
- Non-Workflow mode without Ticket Management—An activity is created automatically and transparently for you whenever you define, modify, or assign policies to devices. The same activity is used until you submit your changes to the database, and is automatically closed and reopened as needed. You cannot actively open or manage activities in non-Workflow mode. These types of activities are also called configuration sessions.
- Workflow mode—If you do not explicitly open an activity, you are prompted to create a new activity or open an existing one whenever you perform an action that requires an activity. You must actively open and manage activities in Workflow mode.

Note Workflow mode works in the same manner whether Ticket Management is enabled or not. Enabling Ticket Management in Workflow mode simply enables the Ticket field for use with Activities. Entering a ticket ID is not required, but if one is used, the Ticket field can be configured to link to an external change management system. For more information, see Ticket Management.

When you create an activity, or one is created for you, you open a virtual copy of the Security Manager policy database. You define and assign policies within this copy. Changes that you made within this copy are only available within the copy. Other users in different activities cannot see these changes. After the activity is submitted and, in Workflow mode, approved, the changes within this copy are committed to the database so that all other users can view the changes. Then, you can create a deployment job to generate the relevant CLI commands and deploy them to the devices.

How you submit your activity changes differs depending on Workflow mode:

- Non-Workflow mode with Ticket Management (default)—Select **Tickets > Submit Ticket** to submit your changes to the policy database.
- Non-Workflow mode without Ticket Management—Select **File > Submit** to submit your changes to the policy database.
- Workflow mode—Select **Activities > Submit Activity** if you are working with an activity approver, or **Activities > Approve Activity** if you do not have a separate activity approver.

The following topics describe why activities are important and how they operate in Workflow mode:

- Benefits of Activities, page 4-2
- Activity Approval, page 4-3
- Activities and Locking, page 4-3
- Activities and Multiple Users, page 4-4
- Understanding Activity/Ticket States, page 4-4

Benefits of Activities

You use activities to control changes made to policies and policy assignments. Although how activities are implemented depends on the workflow settings you choose, all activities provide the following benefits:

- Audit trail—Activities track changes that are made in Security Manager. You can use this information to determine what changes were made and who made the changes as described in Viewing Activity/Ticket Status and History, page 4-23. For both Workflow and non-Workflow mode, there is also an audit report that provides visibility into activities and other actions, as described in Working with Audit Reports, page 10-19.
- Safety mechanism—Activities provide a means for experimenting with changes. Because you are making the changes to a private database view, if you do not want to implement the changes, you can easily discard the activity or configuration session. For more information, see Discarding an Activity/Ticket, page 4-22.
- Task isolation—The policies that are modified within an activity (or configuration session) are locked from being modified within other activities. This prevents conflicting changes that could make a policy unstable. For more information, see Activities and Locking, page 4-3.

In addition, the changes you make within an activity are visible *only* within the activity. Other users see only the last approved committed configurations, unless they view your activity before you close it (in Workflow mode).

Activity Approval

When you enable Workflow mode, you can choose to operate with or without an activity approver.

If your organization requires a different person with higher permissions to approve activities, you can enable workflow with an approver. When using Workflow mode with an approver, the activity must be approved by a person with the appropriate permissions so the policies can be committed to the database. This approval process at the policy definition level helps to ensure that no inappropriate configurations reach the network devices.

If you choose to operate without an approver, the person defining the policies has the permissions to approve them.

For information about enabling or disabling activity approval and changing the default activity approver, see Workflow Page, page 11-75.

Activities and Locking

To prevent multiple users from making conflicting changes, Security Manager obtains activity-level locks when a user performs certain actions within an activity or configuration session in Workflow or non-Workflow mode. This prevents two or more people from making changes to the same feature policy, policy assignment, or object at the same time.

Security Manager also uses locking to ensure that operations related to the committed configuration always run exclusive of one another. These operations can be divided into two categories:

Operations that change the committed configuration:

- Activity approval, which includes configuration session submission in non-Workflow mode.
- Device deletion.
- Editing device properties.

Operations that read the committed configuration:

- Configuration preview.
- Deployment (in non-Workflow mode).
- Creation of deployment job (in Workflow mode).
- Activity or configuration session validation.

If you are performing an operation that changes the committed configuration, no one can perform any of the operations in either list until this operation is complete. An error message is displayed to the user who tries, indicating the action and activity (or user, in non-Workflow mode) that has the lock. For example, if you are approving an activity (which occurs automatically when an activity is submitted in non-Workflow mode), no one else can delete a device or validate a different activity until the approval is complete. This type of locking is particularly important in multi-user settings as it prevents multiple users from simultaneously making changes to the committed configuration.

If you are performing an operation that reads the committed configuration, no one can perform an operation that changes the committed configuration. For example, if you are validating an activity, another user cannot approve an activity. However, other users can perform another operation that reads

the configuration. For example, if you are validating an activity, another user can create a deployment job. Similarly, if you are previewing the configuration before deployment, another user is permitted to do the same. This is because these two operations are limited to reading the committed configuration; they do not make any changes to it.

Tip Activity locking is broader in scope than policy locking, which is described in Understanding Policy Locking, page 5-8. Policy locking prevents two users from changing the same policy on the same device simultaneously.

Related Topics

- Approving or Rejecting an Activity (Workflow Mode), page 4-21
- Deleting Devices from the Security Manager Inventory, page 3-58
- Viewing or Changing Device Properties, page 3-41
- Working with Deployment and the Configuration Archive, page 8-25
- Validating an Activity/Ticket, page 4-18

Activities and Multiple Users

Only one user can define or change policies within an individual activity at one time. However, when Workflow mode is enabled or when Ticket Management is enabled in non-Workflow mode, multiple users can work in the activity in sequence. That is, if an activity or ticket is closed (but not yet approved or submitted for approval), another user can open it and make changes to it if they have the necessary privileges. Multiple users can work in parallel in different activities.

Understanding Activity/Ticket States

Activities in Workflow mode and tickets in non-Workflow mode (when Ticket Management is enabled) can have the states described in the following table. The main states are shown in bold.

Table 4-1 Activity/Ticket States

State	Description
Edit	The activity/ticket was created, but the activity is not currently being edited. The activity/ticket can be opened or discarded while it is in the Edit state.
Edit Open	The activity/ticket is open for editing. Changes, such as defining and assigning policies, can be made in the activity/ticket. The policies, policy assignments (devices being assigned policies), and objects being configured or modified in the activity/ticket are locked. That is, they cannot be configured or modified within the context of another activity/ticket. An activity can be closed, discarded, submitted, or approved while it is in the Edit Open state. A ticket can be closed, discarded, or submitted while it is in the Edit Open state. The configuration changes can be seen only in the context of the activity/ticket.

Table 4-1 Activity/Ticket States (continued)

State	Description
Submitted Submitted Open	The activity was submitted for review and approval or the ticket was submitted. (In Workflow mode, this state is available only if you have activity approval required. For more information, see Workflow Page, page 11-75.) No further changes can be made within the activity/ticket. The policies, devices (through policy assignment), or objects affected by the policy changes remain locked to other activities/tickets. When an activity is submitted, an e-mail is sent to the approver. The approver can open the activity (in read-only mode, moving to the Submitted Open state) to review the changes within the activity, then approve or reject it. An approved activity moves to the approved state. A rejected activity returns to the Edit state.
Approved (Workflow mode only)	The activity was approved, and the corresponding configuration elements are now committed policy configurations. The devices affected by the policy changes are no longer locked to other activities. The activity can be deployed while it is in the Approved state.
Approve Failed (Workflow mode only)	The activity is placed in the Approve Failed state if errors occur during approval (for example, due to a power failure). If this happens, try to approve the activity again or reboot the server.
Discarded	Changes made to the activity/ticket since it was created were discarded and further changes to the activity/ticket are not allowed. Devices associated with the activity/ticket are unlocked and can now be used in a new activity/ticket. The activity/ticket remains in the table showing a Discarded state until it is purged from the system.

Figure 4-1 shows the stages in the ticket workflow. Figure 4-2 shows the stages in the activity workflow without an approver. Figure 4-3 shows the stages in the activity workflow with an approver.

Figure 4-1 Ticket Workflow

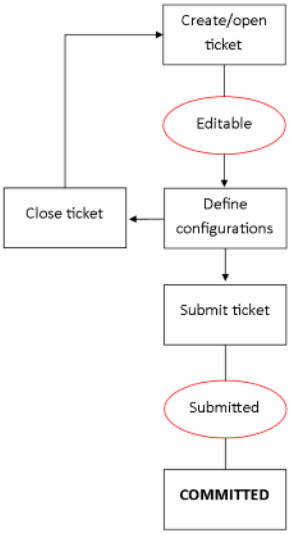

Figure 4-2 Activity Workflow without an Approver

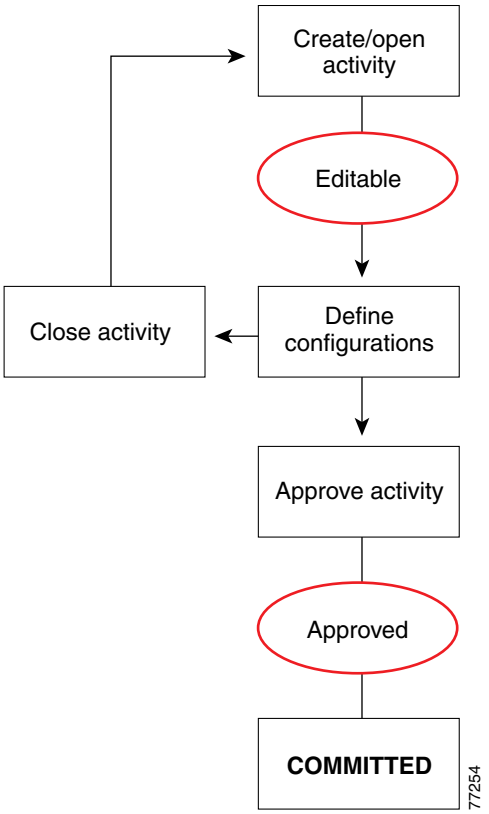

Figure 4-3 Activity Workflow with an Approver

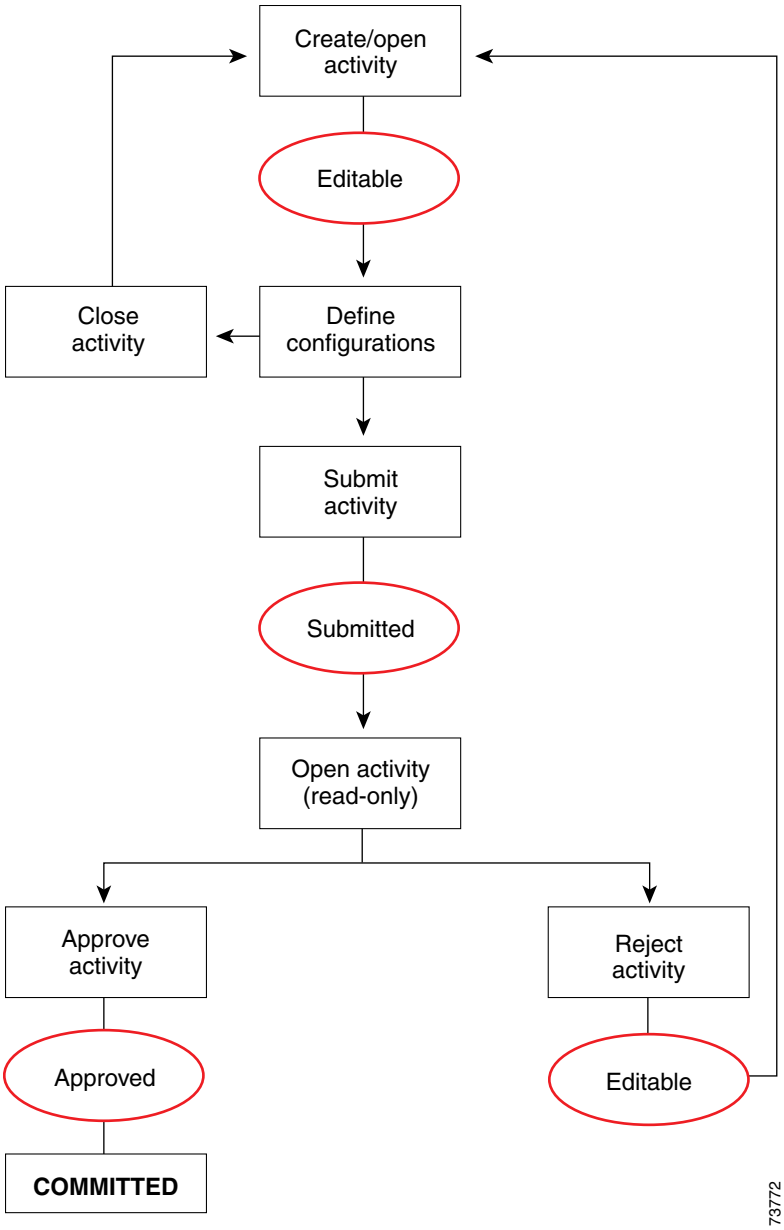

Working with Activities/Tickets

The following topics provide information to help you use activities and configuration sessions:

- Accessing Activity Functions in Workflow Mode, page 4-8
- Accessing Ticket Functions in Non-Workflow Mode, page 4-9
- Creating an Activity/Ticket, page 4-14
- Opening an Activity/Ticket, page 4-15

- Closing an Activity/Ticket, page 4-16
- Viewing Change Reports, page 4-16
- Validating an Activity/Ticket, page 4-18
- Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20
- Approving or Rejecting an Activity (Workflow Mode), page 4-21
- Discarding an Activity/Ticket, page 4-22
- Viewing Activity/Ticket Status and History, page 4-23

Accessing Activity Functions in Workflow Mode

In Workflow mode, you can access activity management functions in the following ways:

- Select **Manage > Activities**. The Activity Manager window contains a list of existing activities and their states. From this window, you can create new activities, and open, close, submit, approve, reject, or discard existing activities. For more information, see Activity/Ticket Manager Window, page 4-10.
- Click a button in the Activities portion of the main toolbar or select the equivalent command in the Activities menu. Whether a button or command is active depends on your user permissions, the state of the activity, and whether you are using workflow with or without an approver. The following table explains the buttons and commands and the conditions under which you can them.

Note If an activity is open, the activity name is displayed above the Global Search field in the upper right corner of the Configuration Manager interface. You can click the activity name to open the Activity Manager window.

Table 4-2 Activities Tool Bar Buttons and Commands When Workflow Mode Is Enabled

Button	Activities Menu Command	Description
	New Activity	Creates an activity.
	Open Activity	Opens an activity. You can open an activity when it is in the Edit or the Submitted state. To open a submitted activity, you must have user privileges to approve or reject changes made in that activity. For more information, see the *Installation Guide for Cisco Security Manager*.
	Close Activity	Saves all changes made while the activity was open and closes it. You can close an activity when it is in the Edit Open or the Submit Open state.
	View Changes	Evaluates all changes made in the activity and produces an Activity Change Report in PDF format in a separate window. For more information, see Viewing Change Reports, page 4-16
	Validate Activity	Validates the integrity of changed policies within the current activity. By validating an activity, you can check for configuration errors that you might have introduced by your policy changes.

Table 4-2 Activities Tool Bar Buttons and Commands When Workflow Mode Is Enabled

Button	Activities Menu Command	Description
	Submit Activity	In Workflow mode with an activity approver, submits the activity for approval. You can submit an activity when it is in the Edit or the Edit Open state.
	Approve Activity	Approves the changes proposed in an activity. You can approve an activity when it is in the Submitted state when using an activity approver, or the Edit or Edit Open state when not using an approver. You must have user privileges to accept the changes proposed in an activity. For more information, see the *Installation Guide for Cisco Security Manager*.
	Reject Activity	In Workflow mode with an activity approver, rejects the changes proposed in an activity. You can reject an activity when it is in the Submitted or Submitted Open state. You must have user privileges to deny changes proposed in an activity. For more information, see the *Installation Guide for Cisco Security Manager*.
	Discard Activity	Discards the selected activity. The activity is discarded and later purged from the system after it exceeds the age for keeping activities as set under Tools > Security Manager Administration > Workflow. The activity state is shown as discarded until the activity is actually purged from the system.

Accessing Ticket Functions in Non-Workflow Mode

In non-Workflow mode with Ticket Management enabled, you can access ticket management functions in the following ways:

- Select **Manage > Tickets**. The Ticket Manager window contains a list of existing tickets and their states. From this window, you can create new tickets, and open, close, submit, or discard existing tickets. For more information, see Activity/Ticket Manager Window, page 4-10.
- Click a button in the Tickets portion of the main toolbar or select the equivalent command in the Tickets menu. Whether a button or command is active depends on your user permissions and the state of the ticket. The following table explains the buttons and commands and the conditions under which you can them.

Note If a ticket is open, the ticket ID is displayed above the Global Search field in the upper right corner of the Configuration Manager interface. You can click the ticket ID to open the Ticket Manager window.

Table 4-3 Tickets Tool Bar Buttons and Commands When Ticket Management Is Enabled in Non-Workflow Mode

Button	Activities Menu Command	Description
	New Ticket	Creates a ticket.
	Open Ticket	Opens a ticket. You can open a ticket when it is in the Edit state.
	Close Ticket	Saves all changes made while the ticket was open and closes it. You can close a ticket when it is in the Edit Open state.
	View Changes	Evaluates all changes made in the ticket and produces a Ticket Change Report in PDF format in a separate window. For more information, see Viewing Change Reports, page 4-16
	Validate Ticket	Validates the integrity of changed policies within the current ticket. By validating a ticket, you can check for configuration errors that you might have introduced by your policy changes.
	Submit Ticket	Submits the ticket. Submitting the ticket saves the proposed changes to the database. Devices associated with the ticket are unlocked, meaning they can be included in policy definitions and changes in other tickets. You can submit a ticket when it is in the Edit or the Edit Open state.
	Discard Ticket	Discards the selected ticket. The ticket is discarded and later purged from the system after it exceeds the age for keeping tickets as set under Tools > Security Manager Administration > Ticket Management. The ticket state is shown as discarded until the ticket is actually purged from the system. For more information, see Ticket Management Page, page 11-72.

Activity/Ticket Manager Window

Activity management and ticket management are very similar processes. The primary difference between activities and tickets is that tickets do not use an approval process. For a comparison of the various modes of operation, see Comparing Workflow Modes, page 1-22.

- **Activity Manager**—Use the Activity Manager window to create and manage activities and to view activity status and history. The upper pane lists the activities that have been created. Select an activity to view its details and history in the lower pane.
- **Ticket Manager**—Use the Ticket Manager window to create and manage tickets and to view ticket status and history. The upper pane lists the tickets that have been created. Select a ticket to view its details and history in the lower pane.

Note The Activity Manager window is available only if you are operating in Workflow mode. The Ticket Manager window is available only if you are operating in non-Workflow mode with Ticket Management enabled. In non-Workflow mode without Ticket Management enabled, Security Manager automatically and transparently manages activities. For information on selecting a mode, see Changing Workflow Modes, page 1-28.

Navigation Path

- In non-Workflow mode with Ticket Management enabled, click the Ticket Manager button on the Main toolbar, or select **Manage > Tickets**.
- In Workflow mode, click the Activity Manager button on the Main toolbar, or select **Manage > Activities**.

Field Reference

Table 4-4 Activity/Ticket Manager Window

Element	Description
*	Activities with unapproved changes (Edit, Edit Open, or Submitted state) or tickets with unsubmitted changes (Edit or Edit Open state) are flagged for easy identification.
Activity Ticket	The name of the activity or ID of the ticket. If Ticket Management is enabled in Workflow mode, both columns are displayed. If Ticket Management is enabled, you can click the ticket ID to view details of the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Last Modified	The date and time of the most recent change to the activity/ticket.
State	The state of the activity/ticket. For a list of states, see Understanding Activity/Ticket States, page 4-4.
User	The username of the person who last changed the state of the activity/ticket.
Last Action	The most recent action performed on the activity/ticket.
Create button	Click this button to create a new activity or ticket so that you can create or change policies or assign policies to devices. For more information, see Creating an Activity/Ticket, page 4-14.
Open button	Click this button to open the selected activity/ticket so that changes, such as defining and assigning policies, are captured within the activity/ticket. You can open an activity when it is in the Edit or the Submitted state. Submitted activities are opened read-only. You can open a ticket when it is in the Edit state. For more information, see Opening an Activity/Ticket, page 4-15.
Close button	Click this button to close the selected activity/ticket if you or others want to continue configuring policies at a later time. For more information, see Closing an Activity/Ticket, page 4-16.
Validate button	Click this button to validate changes that you have made to the selected activity/ticket from the time you created it to the current time. Validating an activity/ticket checks policy integrity and deployability, and displays detailed error information if errors are detected. For more information, see Validating an Activity/Ticket, page 4-18.

Table 4-4 Activity/Ticket Manager Window (continued)

Element	Description
Submit button	In Workflow mode with an activity approver, click this button to submit the selected activity. Submitting the activity sends notification that the activity is ready for review to the specified approver. You can submit an activity when it is in the Edit or the Edit Open state. In non-Workflow mode with Ticket Management enabled, click this button to submit the selected ticket. Submitting the ticket saves the proposed changes to the database. Devices associated with the ticket are unlocked, meaning they can be included in policy definitions and changes in other tickets. You can submit a ticket when it is in the Edit or the Edit Open state. You are prompted for a comment. For more information, see Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20.
Approve button (Activity Manager only)	Click this button to approve the selected activity, which saves the proposed changes to the database. Devices associated with the activity are unlocked, meaning they can be included in policy definitions and changes in other activities. You must have appropriate user permissions to approve the activity. In Workflow mode without an approver, you can approve your own activities when they are in the Edit state. In workflow mode with an approver, you must submit your activity, and the approver can approve an activity only when it is in the Submitted state. You are prompted for an approval comment. For more information, see Approving or Rejecting an Activity (Workflow Mode), page 4-21.
Reject button (Activity Manager only)	In Workflow mode with an activity approver, click this button to reject the changes proposed in the selected activity or activities. You must have appropriate user permissions to reject an activity. If the activity is rejected, the submitter can continue to make changes to the activity. Devices associated with the activity are not unlocked, meaning that they cannot be included in policy definitions or changes in another activity. You can reject an activity only when it is in the Submitted or the Submitted Open state. You are prompted for a rejection comment. For more information, see Approving or Rejecting an Activity (Workflow Mode), page 4-21.

Table 4-4 Activity/Ticket Manager Window (continued)

Element	Description
Discard button	Click this button to discard the selected activity/ticket. Devices associated with the activity/ticket are unlocked, meaning they can be used by other activities/tickets. Multiple activities/tickets can be discarded at the same time. You are prompted for a comment. For more information, see Discarding an Activity/Ticket, page 4-22. Discarded activities are removed from the system according to the settings defined in the Security Manager settings for Workflow. The activity state is shown as discarded until the activity is purged from the system. For more information, see Workflow Page, page 11-75. Discarded tickets are removed from the system according to the settings defined in the Security Manager settings for Ticket Management. The ticket state is shown as discarded until the ticket is purged from the system. For more information, see Ticket Management Page, page 11-72.
View Changes	Click this button to generate a report in PDF format for the selected activity/ticket. If the activity/ticket is closed, this button is grayed out. For more information, see Viewing Change Reports, page 4-16.
Refresh button	Click this button to refresh the information presented in the window.
Details tab	Displays detailed information for the selected activity/ticket. Besides the information repeated from the table, the details include this information: • Activity ID—The identification number assigned by Security Manager when you created the activity. • Ticket ID—The identification number entered when the ticket was created. You can click the Edit Ticket button next to the ticket ID to edit the ticket ID. • Created—The date and time the activity/ticket was created. • Last Modified—The date and time changes were last made to the activity/ticket. • Description—The description that was entered when the activity/ticket was created. • Comments History—Shows a history of the comments that were entered for this activity/ticket. The user that entered the comment is shown as well as the date and time the comment was entered. You can add and edit comments using the buttons below the Comments History table.
History tab	Displays a log of the changes that have been made to the selected activity/ticket. The information includes the state changes, the user who made the change, the date and time of the change (based on the Security Manager server time), and any comments the user entered to document the change.

Creating an Activity/Ticket

In Workflow mode, before you create or change policies or assign policies to devices, you must create an activity. In non-Workflow mode, if you have Ticket Management enabled, before you create or change policies or assign policies to devices, you must create a ticket.

Tip In non-Workflow mode with Ticket Management disabled, activities are created automatically when needed.

Related Topics

- Understanding Activities, page 4-1
- Opening an Activity/Ticket, page 4-15

Step 1 Do one of the following:

For activities:

- Click the **Create Activity** button in the activity toolbar.
- Select **Activities > New Activity**.
- Click **Create** in the Activity Manager window.

For tickets:

- Click the **Create Ticket** button in the tickets toolbar.
- Select **Tickets > New Ticket**.
- Click **Create** in the Ticket Manager window.

The Create Activity/Ticket dialog box appears.

Step 2 In the Create Activity/Ticket dialog box, enter a name for the activity or keep the system-generated name. The default name contains the username, date, and time the activity/ticket was created. You can also enter a comment to describe the activity/ticket.

Ticket Management supports linking between a Ticket ID and an external ticket management system. For more information, see Ticket Management Page, page 11-72.

Tip You can use a comma to separate multiple ticket IDs.

Step 3 Click **OK**.

The activity/ticket is listed in the Activity/Ticket Manager window. For more information, see Activity/Ticket Manager Window, page 4-10.

Responding to the Activity/Ticket Required Dialog Box

When in Workflow mode, you must create or open an activity before you create or modify policies. When in non-Workflow mode, if Ticket Management is enabled, you must create or open a ticket before you create or modify policies. If you attempt to perform an action that requires an activity or a ticket, and you have not created or opened one yet, you are prompted to do so with the Activity/Ticket Required dialog box.

You can choose from the following options:

- **Create a new activity/ticket**—Create a completely new activity/ticket, specifying an activity name or ticket ID and optionally a description of the purpose of the activity/ticket. The default activity/ticket name contains the username, date, and time the activity/ticket was created.
- **Open an existing activity/ticket**—To open the activity/ticket you select from the Activity/Ticket list. This option is displayed only if there are activities/tickets available in the Edit state.

Related Topics

- Creating an Activity/Ticket, page 4-14
- Understanding Activity/Ticket States, page 4-4

Opening an Activity/Ticket

In Workflow mode, you can open an existing activity if no one else has it opened. You might open an existing activity in the Edit state to make further policy changes, or you might open an existing activity in the Submitted state to review proposed policy changes before approving or rejecting it (if you have the appropriate permissions and you are working in Workflow mode with an approver).

You can make changes to activities in the Edit state, but you can only view activities in the Submitted state.

In non-Workflow mode, if you have Ticket Management enabled, you can open an existing ticket to make further policy changes if no one else has it opened.

To open an activity/ticket, do one of the following:

- For activities:
  - Click the **Open** button in the activity toolbar or select **Activities > Open Activity**. The Openable Activities dialog box lists all activities that can be opened, including the name of the activity, its state, and the username of the person who created the activity. Select the activity you want to open and click **OK**.
  - Select **Manage > Activities**. From the Activity Manager window, select the activity you want to open and click **Open**.
- For tickets:
  - Click the **Open** button in the tickets toolbar or select **Tickets > Open Ticket**. The Openable Tickets dialog box lists all tickets that can be opened, including the ticket ID, its state, and the username of the person who created the ticket. Select the ticket you want to open and click **OK**.
  - Select **Manage > Tickets**. From the Ticket Manager window, select the ticket you want to open and click **Open**.

Tip In non-Workflow mode with Ticket Management disabled, your previous configuration session is opened whenever needed until you submit it. A new activity is then created the next time you perform an action that requires an activity.

Note In Workflow mode and in non-Workflow mode with Ticket Management enabled, you are prompted to open or create an activity/ticket when you launch Security Manager.

Related Topics

- Understanding Activities, page 4-1

Closing an Activity/Ticket

You can close an activity without approving it (or submitting it for approval) or close a ticket without submitting it if you or others want to continue configuring policies at a later time.

A person with administrator privileges can close an activity/ticket opened by another user.

To close an open activity/ticket, do one of the following:

- For activities:
  - Click the **Close** button in the activity toolbar.
  - Select **Activities > Close Activity**.
  - Select **Manage > Activities**. From the Activity Manager window, click **Close**.
- For tickets:
  - Click the **Close** button in the tickets toolbar.
  - Select **Tickets > Close Ticket**.
  - Select **Manage > Tickets**. From the Ticket Manager window, click **Close**.

Tip In non-Workflow mode with Ticket Management disabled, your configuration session is closed whenever you log out. The same session is reopened the next time you log in.

Related Topics

- Understanding Activities, page 4-1

Viewing Change Reports

There are many places in the interface where you can open change reports. Typically, the button or command to generate the report is **View Changes**. These change reports provide detailed information about the policy and policy object changes, and the devices that were acted on, that have been made in an activity/ticket, whether you are operating in Workflow or non-Workflow mode.

The change report is in Adobe Acrobat (PDF) format. You can use all of the Acrobat features, including the bookmarks tab, to view the report.

If you discover a device or rediscover policies on a device, then subsequent policy changes in the same activity/ticket performed on that device are not listed in the activity change report. This is also true on a device that you clone from another device.

Following are some of the ways you can view change reports:

- Non-Workflow mode with Ticket Management:
  - Select **Tickets > View Changes**, or click the **View Changes** button in the toolbar, to view the changes made during the currently open ticket.
  - Highlight a ticket in the Ticket Manager window and click **View Changes** to view the changes made in that ticket
- Non-Workflow mode without Ticket Management:

  - Select **File > View Changes** to view the changes made during the current configuration session.
  - Select **Manage > Change Reports** to view the changes made during previous sessions (which are closed when you submit or discard your changes). Select a configuration session from the Change Report window and click **View Changes**. (See Selecting a Change Report in Non-Workflow Mode with Ticket Management Disabled, page 4-18.)
- Workflow mode:
  - Select **Activities > View Changes**, or click the **View Changes** button in the toolbar, to view the changes made during the currently open activity.
  - Highlight an activity in the Activity Manager window and click **View Changes** to view the changes made in that activity.
- In all modes, you can view changes from various dialog boxes when creating deployment jobs.

Note You must disable any popup-blocker applications you have running to ensure the activity report will open.

The following illustration shows a sample activity report.

Figure 4-4 Activity Change Report

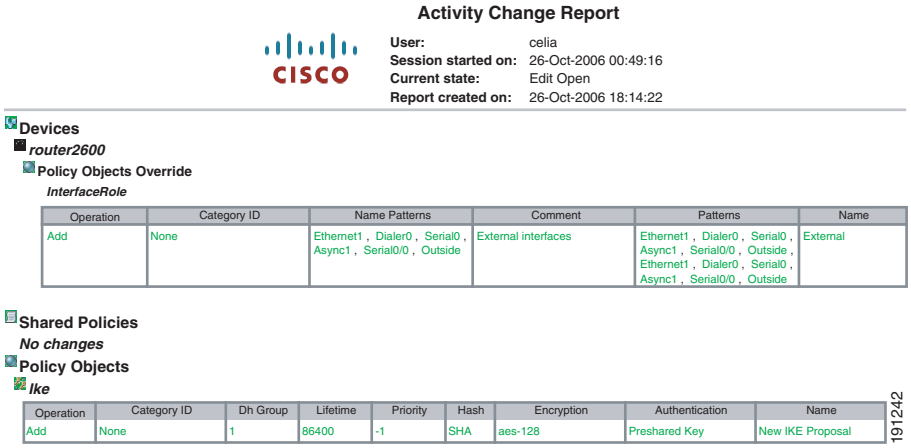

The change report includes these elements:

- Activity name/Ticket ID—The name of the activity (or the user and session start date and time if it is unnamed) or the Ticket ID.
- Created by—The username of the person who created the activity/ticket, with the date and time.
- Current state—The current state of the activity.
- Report created on—The date and time the report was created.
- Devices section—A summary of the devices that were acted on in the activity/ticket (that is, they were added, modified, or deleted). Changes to local policies are displayed here.

  Changes in this section and the other sections of the report are color-coded to help you identify changes:
  - Green—Indicates a newly inserted item.
  - Red—Indicates a deleted item or the old value of a changed item.

  - Blue—Indicates the new value of a changed item.
- Shared Policies section—Changes to all shared policies are displayed here.
- Policy Bundles—Changes to all policy bundles are displayed here.
- Policy Objects—Changes to all policy objects are displayed here.
- VPN—Changes to VPN topologies and policies are displayed here, including newly discovered VPNs and deleted VPN topologies.

Selecting a Change Report in Non-Workflow Mode with Ticket Management Disabled

In non-Workflow mode with Ticket Management disabled, you can view change reports for closed configuration sessions by selecting **Manage > Change Reports** and then selecting the session in the Change Report dialog box.

In non-Workflow mode with Ticket Management disabled, a configuration session is considered complete when you either submit or discard your changes. The Change Report dialog box lists all closed sessions, showing the date and time the session was closed, the action that closed it (submitted or discarded), and the user name associated with the session. These sessions are equivalent to activities in Workflow mode. Select a session and click **View Changes** to view the report. For information on reading the report, see Viewing Change Reports, page 4-16.

Tip To view the report for the current configuration session, close this dialog box and select **File > View Changes**.

Validating an Activity/Ticket

In Workflow mode, Security Manager validates activities when you submit them for approval, or you can validate an activity at any time while you are creating and changing policies in an activity. After an activity is submitted, the validation report remains static.

In non-Workflow mode, Security Manager validates policies when you submit them to the database, when you try to deploy them, or when you validate them. The validation process reports on policy changes that were made up until the changes are submitted or deployed.

The validation process checks the following areas. If there are errors, you can display a detailed summary of the validation results.

- Policy integrity—There are no unresolvable references (for example, missing objects, unresolved interface roles, overrides of mandatory settings, and so on).
- Policy deployability—The platform, operating system, and configured features are supported by the target devices so that policies can be correctly translated into CLI commands.

  If a policy contains options that require specific device types or operation system versions, you will see validation warnings for non-supported devices, but Security Manager will not generate the associated commands for unsupported devices. This allows you to create policies that apply to a wide range of devices without having to create policies that are too device-specific.
- FlexConfig integrity—There are no corrupted FlexConfig objects. If corrupted objects are found, a warning with a list of the corrupted FlexConfig objects results.
- FlexConfig syntax—If syntax errors are found, a warning with a list of affected FlexConfigs and their syntax errors results.

- FlexConfig object references—All object references are resolvable. If FlexConfig objects reference non-existent objects, a warning with a list of the missing objects results.

Related Topics

- Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20
- Deploying Configurations in Non-Workflow Mode, page 8-28

Step 1 Do one of the following:

- In Workflow mode:
  - Open an activity, and then click the **Validate** button on the activity toolbar or select **Activities > Validate Activity**.
  - Select **Manage > Activities**. From the Activity Manager window, select an activity, and then click **Validate**.
- In non-Workflow mode with Ticket Management enabled:
  - Open a ticket, and then click the **Validate Saved Changes** button on the tickets toolbar or select **Tickets > Validate Ticket**.
  - Select **Manage > Tickets**. From the Ticket Manager window, select a ticket, and then click **Validate**.
- In non-Workflow mode with Ticket Management disabled, select **File > Validate**, or try to preview or deploy policies.

Security Manager performs the validation and opens an informational message dialog box that summarizes the validation results. If there are no errors, validation passes. If there are errors or warnings, click **Details** to open the Validation dialog box, where you can view detailed information about the errors.

Step 2 Evaluate the errors to determine how to fix them.

The Validation dialog box organizes the errors and warnings in two ways, which are displayed on separate tabs:

- Errors tab—The Errors tab organizes validation problems based on the type of error. Each error indicates the number of devices that are affected and the severity of the error.

  Select an error in the upper pane, and a list of devices (with the type of device) that have the error appears in the lower left pane. The lower right pane describes the error, its cause, and what you might do to fix it.
- Devices tab—The Devices tab organizes validation problems based on the device. Each device indicates the number and types of errors and warnings for the device, and the device type. The device status indicates the worst problem in the device configuration (error or warning).

  Select a device in the upper pane, and a list of the errors for that device appears in the lower left pane. Select an error and the lower right pane describes the error, its cause, and what you might do to fix it.

You must correct errors before submitting the activity. Security Manager does not allow an activity to be submitted with validation errors.

Note A validation warning (as opposed to an error) will not prevent activity approval or deployment.

Submitting an Activity for Approval (Workflow Mode with Activity Approver)

In Workflow mode with an activity approver, you must submit activities for approval. When you submit the activity, the integrity and deployability of the activity is validated. For details about the validation process and report, see Validating an Activity/Ticket, page 4-18.

The activity is also closed so that it can be opened by the user who has the permissions to approve it. When the activity is approved, its configurations are committed to the Security Manager database, and they can be deployed to the devices.

When you submit an activity, Security Manager sends an e-mail to the relevant approvers to notify them that an activity requires approval.

If you are working in Workflow mode without an activity approver, you do not need to submit activities (in fact, you cannot submit them). You can approve the activity yourself. For more information about changing activity approval settings, and configuring the e-mail addresses for notifications, see Workflow Page, page 11-75.

Related Topics

- Understanding Activities, page 4-1
- Opening an Activity/Ticket, page 4-15
- Understanding Activity/Ticket States, page 4-4
- Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27

Step 1 Do one of the following:

- Open an activity and click the **Submit Activity** button on the activity toolbar or select **Activities > Submit Activity**.
- Select **Manage > Activities**. From the Activity Manager window, select an activity, then click **Submit**.

The Submit Activity dialog box opens.

Step 2 In the Submit Activity dialog box, fill in the following fields:

- **Approver**—Enter the e-mail address of the person who should approve the activity if the default address is not the right one. This person receives notification of your submission.

  The default e-mail address is set in Tools > Security Manager Administration > Workflow.
- **Comment**—Enter comments that will help the approver evaluate the activity.
- **Submitter**—Enter the e-mail address of the person submitting the approval request if the default address is not the right one. The field initially contains the e-mail address associated with the username you used to log into Security Manager. Notifications of activity state changes are sent to this address.

If desired, you can click the **View Changes** button to view a report in PDF format of the changes made in the activity. For more information, see Viewing Change Reports, page 4-16.

Step 3 Click **OK**. The activity status changes to Submitted in the Activity Manager window and notifications are sent.

Note Security Manager warns you if the e-mail cannot be sent and you must contact the approver directly.

Approving or Rejecting an Activity (Workflow Mode)

Before the changes in an activity are committed to the database, you must approve the activity. If you have activity approval permissions, you can open an activity, review the policies and policy assignments, and then either approve or reject the activity.

If you are operating in Workflow mode without an approver, you can approve your own activities. When working without an approver, you cannot reject an activity, but you can discard it if you do not want to save your changes. In non-Workflow mode, you use the Submit and Discard commands on the file menu to submit (and automatically approve) or discard the configuration session.

In Workflow mode with an activity approver, the activity must be submitted before you can open it and approve it. In this mode, you can also reject the activity.

If you approve an activity, policies and policy assignments are committed to the database and are ready to be deployed to devices or files. Devices associated with the activity are unlocked, meaning they can be included in policy definitions and changes in other activities.

If you reject the activity, it is returned to the Edit state and the submitter can reopen the activity to make the necessary changes and resubmit it for approval. Devices associated with the activity are not unlocked, meaning that they cannot be included in policy definitions or changes in another activity.

Note After an activity is approved, changes cannot be undone. You must create a new activity and manually change policies and policy assignments to the desired state.

Related Topics

- Understanding Activities, page 4-1
- Opening an Activity/Ticket, page 4-15
- Understanding Activity/Ticket States, page 4-4

Step 1 Perform the appropriate step below:

- To approve an activity, do one of the following:
  - Open an activity and click the **Approve Activity** button on the activity toolbar or select **Activities > Approve Activity** from the menu.
  - Select **Manage > Activities**. In the Activity Manager window, select an activity and click **Approve**.
- To reject an activity or multiple activities, do one of the following:
  - Open an activity and click the **Reject Activity** button on the activity toolbar or select **Activities > Reject Activity** from the menu.
  - Select **Manage > Activities**. In the Activity Manager window, select an activity or multiple activities and click **Reject**.

The Approve Activity, Reject Activity, or Reject Multiple Activities dialog box appears.

Step 2 In the Comment field, enter a brief explanation of why you are approving or rejecting the activity or activities. If you are rejecting, you might want to include suggested revisions.

Step 3 Click **OK** (for a single activity) or **Reject** (for multiple activities). The activity status changes to Approved or Edit (if rejected) in the Activity Manager window. For a description of the elements in the window, see Activity/Ticket Manager Window, page 4-10.

Discarding an Activity/Ticket

You can discard an activity/ticket (configuration session in non-Workflow mode) if it is no longer required. When you discard an activity/ticket, you delete all the policies and policy assignments that were defined within the activity. Those policies and policy assignments are not in the database; therefore, they cannot be deployed.

Discarded activities are removed from the system according to the settings defined in the Security Manager settings for Workflow and devices associated with the activity are unlocked, meaning they can be used by other activities. For more information, see Workflow Page, page 11-75.

Discarded tickets are removed from the system according to the settings defined in the Security Manager settings for Ticket Management and devices associated with the ticket are unlocked, meaning they can be used by other tickets. The ticket state is shown as discarded until the ticket is purged from the system. For more information, see Ticket Management Page, page 11-72.

Note When a ticket is discarded, Image Management jobs assigned to that ticket will not be discarded automatically. If required, you must search for pending Image Manager jobs with that ticket ID in Image Manager and delete those jobs.

To discard an activity/ticket:

Workflow Mode—Do one of the following:

- To discard a single activity, do one of the following:
  - Open an activity, then click the **Discard** button on the activity toolbar or select **Activities > Discard Activity**.
  - Select **Manage > Activities**. From the Activity Manager window, select an activity, then click **Discard**. Only an activity in the Edit or Edit Open state can be discarded.

  Using either method, you are prompted with the Discard Activity dialog box, which allows you to enter an optional comment to explain why you are discarding the activity. Enter a comment and click **OK** to discard it.
- To discard multiple activities, select **Manage > Activities**. From the Activity Manager window, select the activities you wish to discard, then click **Discard**. Only activities in the Edit or Edit Open state can be discarded.

  You are prompted with the Discard Multiple Activities dialog box, which allows you to enter an optional comment to explain why you are discarding the activities. If you selected activities belonging to other users, you can choose to discard those activities or not using the Discard selected activities of other users as well check box. Enter a comment, select or deselect the Discard selected activities of other users as well check box as desired, and then click **OK** to discard the selected activities.

Non-Workflow Mode with Ticket Management Enabled—Do one of the following:

- To discard a single ticket, do one of the following:
  - Open a ticket, then click the **Discard** button on the tickets toolbar or select **Tickets > Discard Ticket**.
  - Select **Manage > Tickets**. From the Ticket Manager window, select a ticket, then click **Discard**. Only a ticket in the Edit or Edit Open state can be discarded.

  Using either method, you are prompted with the Discard Ticket dialog box, which allows you to enter an optional comment to explain why you are discarding the ticket. Enter a comment and click **OK** to discard it.
- To discard multiple activities, select **Manage > Tickets**. From the Ticket Manager window, select the tickets you want to discard, then click **Discard**. Only tickets in the Edit or Edit Open state can be discarded.

  You are prompted with the Discard Multiple Tickets dialog box, which allows you to enter an optional comment to explain why you are discarding the tickets. If you selected tickets belonging to other users, you can choose to discard those tickets or not using the Discard selected tickets of other users as well check box. Enter a comment, select or deselect the Discard selected tickets of other users as well check box as desired, and then click **OK** to discard the selected tickets.

Non-Workflow Mode with Ticket Management Disabled—Select **File > Discard** to discard the changes in the current configuration session.

Related Topics

- Understanding Activities, page 4-1
- Opening an Activity/Ticket, page 4-15
- Understanding Activity/Ticket States, page 4-4

Viewing Activity/Ticket Status and History

In Workflow mode, you can view the status and history of changes for activities in the Activity Manager window. In non-Workflow mode with Ticket Management enabled, you can view the status and history of changes for tickets in the Ticket Manager window.

To open the window for activities, click the Activity Manager button in the toolbar or select **Manage > Activities**.

To open the window for tickets, click the Ticket Manager button in the toolbar or select **Manage > Tickets**.

The upper pane lists all available activities/tickets, including the current state of the activity/ticket. Select an activity/ticket to see additional information in the tabs in the lower pane:

- Details tab—Shows the date and time the activity/ticket was created, and its description.
- History tab—Shows the transaction history for the activity/ticket. Each time the state is changed, a record of the change is kept, including the user who made the change and any comments about the change.

Related Topics

- Understanding Activities, page 4-1
- Activity/Ticket Manager Window, page 4-10

Managing Policies

The following topics describe the concept of policies in Cisco Security Manager and how to use and manage them.

- Understanding Policies, page 5-1
- Discovering Policies, page 5-12
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Managing Shared Policies in Policy View, page 5-50
- Managing Policy Bundles, page 5-57

Understanding Policies

In Security Manager, a policy is a set of rules or parameters that define a particular aspect of network configuration. You configure your network by defining policies on devices (which includes individual devices, service modules, security contexts, and virtual sensors) and VPN topologies (which are made up of multiple devices), and then deploying the configurations defined by these policies to these devices.

Several types of policies might be required to configure a particular solution. For example, to configure a site-to-site VPN, you might need to configure multiple policies, such as IPsec, IKE, GRE, and so forth.

Policies are assigned to one or more devices. After a policy is assigned to a device, any changes to the policy definition change the behavior of the device.

The following topics describe policies in more detail:

- Settings-Based Policies vs. Rule-Based Policies, page 5-2
- Service Policies vs. Platform-Specific Policies, page 5-2
- Local Policies vs. Shared Policies, page 5-3
- Understanding Rule Inheritance, page 5-4
- Policy Management and Objects, page 5-7
- Understanding Policy Locking, page 5-8
- Customizing Policy Management for Routers and Firewall Devices, page 5-11

Settings-Based Policies vs. Rule-Based Policies

Security Manager policies are structured as either rule-based policies or settings-based policies.

Rule-Based Policies

Rule-based policies contain one or more rules that govern how to handle traffic on a selected device, such as the access rules and inspection rules defined as part of a firewall service. Rule-based policies can contain hundreds or even thousands of rules arranged in a table, each defining different values for the same set of parameters. The ordering of the rules is very important, as traffic flows are assigned the first rule whose definition matches the flow (known as first matching).

The structure of the rules table depends on whether you configure a local policy or a shared policy (see Local Policies vs. Shared Policies, page 5-3). If you configure a local rule-based policy for a single device, the policy contains a flat table of local rules. If you configure a shared rule-based policy (either in Device view or Policy view), the table is divided into two sections, Mandatory and Default. Mandatory rules always precede the default rules, and cannot be overridden by local or default rules. The Default section contains rules that can be overridden by mandatory and local rules. You can define rules in either the Mandatory or Default section and move rules between sections using cut-and-paste.

When you define certain types of rule-based policies, such as firewall service policies, you can create a policy hierarchy in which rules located at lower levels in the hierarchy acquire properties from the rules located above them. This is known as rule inheritance. For example, you can define a set of inspection rules that apply globally to all firewalls, while supplementing these rules with additional rules that can be applied to a subset of devices. By maintaining common rules in a parent policy, inheritance enables you to reduce the chance of introducing configuration errors that will cause deployment to fail. For more information, see Understanding Rule Inheritance, page 5-4.

Settings-Based Policies

Settings-based policies contain sets of related parameters that together define one aspect of security or device operation. For example, when you configure a Cisco IOS router, you can define a quality of service (QoS) policy that defines which interfaces are included in the policy, the type of traffic on which QoS is applied, and the definition of how this traffic should be queued and shaped. Unlike rule-based policies, which can contain hundreds of rules containing values for the same set of parameters, you can define only one set of parameters for each settings-based policy defined on a device.

Related Topics

- Understanding Policies, page 5-1

Service Policies vs. Platform-Specific Policies

Security Manager policies are divided into several domains, each of which represents a major policy category. These domains can be divided into two categories: service policies and platform-specific policies.

Service policies are divided into the following policy domains:

- Firewall.
- Site-to-site VPN.
- Remote Access VPN.
- IPS service policies.

For example, the firewall policy domain contains policies for access rules, inspection rules, and transparent rules, among others. The site-to-site VPN policy domain contains policies for IKE proposals, IPsec proposals, and preshared keys, among others. Service policies can be applied to any kind of device, regardless of platform, although there may be some variation in policy definition depending on the device type.

Platform-specific policy domains contain policies that configure features that are specific to the selected platform. Not all platform-specific policies are directly related to security. For example, the Router policy domain contains routing policies, identity policies (Network Admission Control and 802.1x), policies related to device administration (DHCP, SNMP, device access), and other policies such as QoS and NAT.

For routers and firewalls (ASA, PIX, FWSM), you can choose which platform-specific policies to manage. For more information, see Customizing Policy Management for Routers and Firewall Devices, page 5-11.

Local Policies vs. Shared Policies

The policies that you configure on devices can either be local or shared. Local policies refer to policies that are defined for a single device. Any changes that you make to a local policy affect only that device. Local policies are well-suited to smaller networks and to devices requiring nonstandard configurations. For example, you might configure a local policy on a router that requires a different OSPF routing policy than the one used by the other routers in your network. For more information about the actions you can perform on local policies, see Performing Basic Policy Management, page 5-31.

As your network grows, maintaining local policies on each device greatly increases the effort required to manage these policies in a comprehensive and efficient manner. To meet this challenge, Security Manager features policy sharing. With policy sharing, you can create a single policy and assign it to multiple devices. For more information, see Sharing a Local Policy, page 5-41.

Figure 5-1 Local vs. Shared Policies

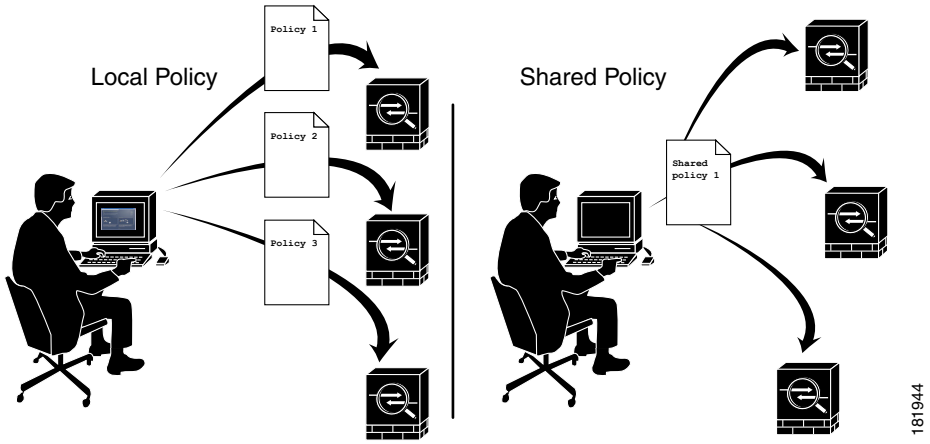

For example, if you want all the Cisco IOS routers in your network to implement the same Network Admission Control (NAC) policy, you need only define a single NAC policy and share it. You can then assign the shared policy to all the routers in your network with a single action. For more information, see Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.

Any changes that you make to a shared policy are automatically applied to all the devices to which it is assigned. As a result, shared policies both streamline the process of policy creation and help maintain consistency and uniformity in your device configurations.

For more information about the actions you can perform on shared policies, see Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37.

Tips

- Shared policies can be grouped together to form policy bundles. Policy bundles make managing the assignment of shared policies easier especially when working with a large number of devices. For more information, see Managing Policy Bundles, page 5-57.
- In addition to sharing policies, you can choose to inherit the rules of a rule-based policy when defining another policy of the same type. This makes it possible, for example, to maintain a set of corporate access rules that apply to all firewall devices while providing the flexibility to define additional rules on individual devices as required. For more information, see Understanding Rule Inheritance, page 5-4.
- If you use more than one Security Manager server, you can maintain a consistent set of policies among the servers by regularly exporting shared policies from your primary server and importing them into the other servers. You must decide which server to use as the official policy source. For more information, see Exporting Shared Policies, page 10-12 and Importing Policies or Devices, page 10-13
- In Version 4.7, Cisco Security Manager has added a new option to the available filtering choices in the Device Filter. This new option provides a filter for devices that have shared policies applied. To see this in the Security Manager GUI, navigate to View > Device View > Filter: > Create Filter... [in the dropdown list]. When the Create Filter dialog box appears, use the dropdown lists to select "Device," "has," and "Shared Policy," for a resulting filter of "Device has 'Shared Policy.'"

Shared Policies and VPNs

In the same way that shared policies facilitate device configuration, they also facilitate the configuration of VPNs. For example, you can create a shared IPsec proposal policy and assign it to multiple site-to-site VPNs. Any changes that you make to the shared policy affect all the VPNs to which the policy is assigned.

You can assign the shared policies to an existing VPN using the Site-to-Site VPN Manager; right-click a shareable policy and select **Assign Shared Policy**. This is done in much the same way as assigning shared policies in Device view. You can also configure shared policies as the default policies to use in the Create VPN wizard, as described in Understanding and Configuring VPN Default Policies, page 25-12.

Related Topics

- Understanding Policies, page 5-1

Understanding Rule Inheritance

As described in Local Policies vs. Shared Policies, page 5-3, shared policies enable you to configure and assign a common policy definition to multiple devices. Rule inheritance takes this feature one step further by enabling a device to contain the rules defined in a shared policy *in addition to* local rules that are specific to that particular device. Using inheritance, Security Manager can enforce a hierarchy where policies at a lower level (called child policies) inherit the rules of policies defined above them in the hierarchy (called parent policies).

Note If a policy bundle includes a shared policy that inherits from other shared policies, those inherited rules are also applied to any devices on which the policy bundle is applied.

Rule Order When Using Inheritance

As described in Understanding Access Rules, page 16-1, an access list (ACL) consists of rules (also called access control entries or ACEs) arranged in a table. An incoming packet is compared against the first rule in the ACL. If the packet matches the rule, the packet is permitted or denied, depending on the rule. If the packet does not match, the packet is compared against the next rule in the table and so forth, until a matching rule is found and executed.

This first-match system means that the order of rules in the table is of critical importance. When you create a shared access rule policy, Security Manager divides the rules table into multiple sections, Mandatory and Default. The Mandatory section contains rules that cannot be overridden by the local rules defined in a child policy. The Default section contains rules that *can* be overridden by local rules.

Figure 5-2 describes how rules are ordered in the rules table when using inheritance.

Figure 5-2 Order of Rules When Using Inheritance

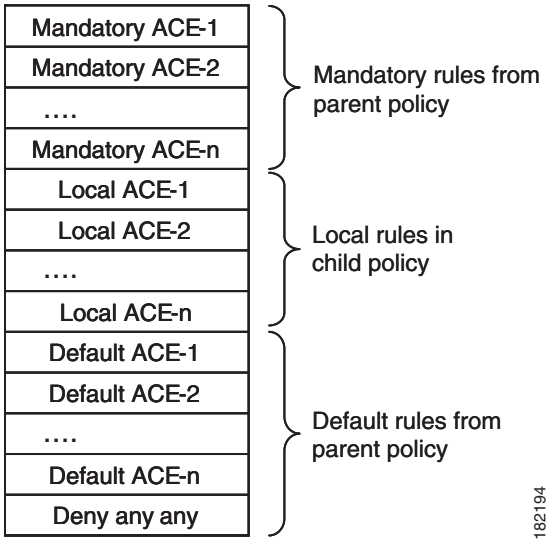

Benefits of Using Inheritance

The ability to define rule-based policies in a hierarchical manner gives you great flexibility when defining your rule sets, and the hierarchy can extend as many levels as required. For example, you can define an access rule policy for the device at a branch office that inherits rules from a parent policy that determines access at the regional level. This policy, in turn, can inherit rules from a global access rules policy at the top of the hierarchy that sets rules at the corporate level.

In this example, the rules are ordered in the rules table as follows:

```
Mandatory corporate access rules
    Mandatory regional access rules
        Local rules on branch device
    Default regional access rules
Default corporate access rules
```

The policy defined on the branch device is a child of the regional policy and a grandchild of the corporate policy. Structuring inheritance in this manner enables you to define mandatory rules at the corporate level that apply to all devices and that cannot be overridden by rules at a lower level in the hierarchy. At the same time, rule inheritance provides the flexibility to add local rules for specific devices where needed.

Having default rules makes it possible to define a global default rule, such as “deny any any”, that appears at the end of all access rule lists and provides a final measure of security should gaps exist in the mandatory rules and default rules that appear above it in the rules table.

Inheritance Example

For example, you can define a mandatory worm mitigation rule in the corporate access rules policy that mitigates or blocks the worm to all devices with a single entry. Devices configured with the regional access rules policy can inherit the worm mitigation rule from the corporate policy while adding rules that apply at the regional level. For example, you can create a rule that allows FTP traffic to all devices in one region while blocking FTP to devices in all other regions. However, the mandatory rule at the corporate level always appears at the top of the access rules list. Any mandatory rules that you define in a child policy are placed after the mandatory rules defined in the parent policy.

With default rules, the order is reversed—default rules defined in a child policy appear before default rules inherited from the parent policy. Default rules appear after any local rules that are defined on the device, which makes it possible to define a local rule that overrides a default rule. For example, if a regional default rule denies FTP traffic to a list of destinations, you can define a local rule that permits one of those destinations.

IPS Policy Inheritance

Event action filter policies for IPS devices can also use inheritance to add rules defined in a parent policy to the local rules defined on a particular device. The only difference is that although active and inactive rules are displayed together in the Security Manager interface, all inactive rules are deployed last, after the inherited default rules.

Signature policies for IPS devices use a different type of inheritance that can be applied on a per-signature basis. See Configuring Signatures, page 39-4.

Related Topics

- Settings-Based Policies vs. Rule-Based Policies, page 5-2
- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Inheritance vs. Assignment, page 5-6
- Inheriting or Uninheriting Rules, page 5-47

Inheritance vs. Assignment

It is important to understand the difference between rule inheritance and policy assignment:

- Inheritance—When you inherit the rules from a selected policy, you do not overwrite the local rules that are already configured on the device. Instead, the inherited rules are *added* to the local rules. If the inherited rules are mandatory rules, they are added before the local rules. If the inherited rules are default rules, they are added *after* the local rules. Any changes that you make to the inherited rules in the parent policy are reflected in the policy that inherits those rules.

Note Inheritance works differently for IPS signature policies and signature event actions. For more information, see Understanding Signature Inheritance, page 39-3.

- Assignment—When you assign a shared policy to a device, you *replace* whatever was already configured on the device with the selected policy. This holds true whether the device previously had a local policy or a different shared policy of that type.

Therefore, when working with rule-based policies such as access rules, you must use discretion when choosing these options. Use inheritance to supplement the local rules on the device with additional rules from a parent policy. Use assignment to replace the policy on the device with a selected shared policy.

Tip To prevent overwriting your local rules by mistake, Security Manager displays a warning message when you select the Assigned Shared Policy option for a rule-based policy. The message provides you the option of inheriting the rules of the policy instead of assigning it. Choose the inheritance option if you want to preserve your local rules.

Related Topics

- Understanding Rule Inheritance, page 5-4
- Inheriting or Uninheriting Rules, page 5-47
- Local Policies vs. Shared Policies, page 5-3
- Settings-Based Policies vs. Rule-Based Policies, page 5-2

Policy Management and Objects

Objects make it easier to configure policies in Security Manager by providing a set of values with a logical, easy-to-remember name that can be applied wherever it is needed. For example, you can define a network/host object called MyNetwork that contains a set of IP addresses in your network. Whenever you configure a policy requiring these addresses, you can simply refer to the MyNetwork object instead of manually entering the addresses each time.

When you define a policy, you can create objects on the fly by clicking the **Select** button next to any field that accepts an object as a value. For more information, see Selecting Objects for Policies, page 6-2. You can also create and manage objects system-wide from the Policy Object Manager, page 6-4.

Policy objects also are created when you discover policies that already exist on a device. You can discover policies when you add a device to the Security Manager inventory, or you can discover policies on devices that are already in the inventory, as described in Discovering Policies, page 5-12. You can configure Security Manager to reuse already-defined policy objects for newly-discovered policies. For more information on configuring policy object settings for discovery, see Discovery Page, page 11-25.

Certain types of objects enable you to override their predefined values at the device level, which enables you to use an object in a policy while retaining the ability to customize particular values. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18.

For more information about objects and how to use them when defining policies, see Chapter 6, “Managing Policy Objects”.

Related Topics

- Understanding Policies, page 5-1

Understanding Policy Locking

Security Manager has a policy locking mechanism that is useful in organizations where several people have the authority to make configuration changes. It prevents a potential situation in which two or more people are making changes to the same device, policy, policy assignment, or object at the same time. When a lock is applied, a message is displayed across the top of the work area to other users who access that device or policy.

Tip Security Manager also obtains activity (or configuration session) locks, which are broader in scope than policy locks, when users perform some actions. For more information, see Activities and Locking, page 4-3.

Lock Types

Security Manager uses two different types of locks:

- Policy content locks—Locks the content of a particular policy. The banner displayed above the work area reads:

  This data for this policy is locked by activity/user: <name>.

  The content lock prevents other users from making any changes to the configuration of the locked policy.

- Assignment locks—Locks the assignment of a policy type to a particular device. The banner displayed above the work area reads:

  The assignment of this policy is locked by activity/user: <name>.

  For a local policy, an assignment lock prevents other users from unassigning the policy or assigning a shared policy of the same type in place of the local policy. For a shared policy, an assignment lock prevents other users from assigning a different shared policy of the same type in place of one already assigned.

These locks can either work together or independently of one another, depending on the actions being performed by the user. If both locks are active at the same time, the banner displayed above the work area reads:

This policy is locked by activity/user: <name>.

See Understanding Locking and Policies, page 5-9 for a summary of the effects locking has on the actions you can perform.

Releasing Locks

After is locked is enabled, it remains in place until you either submit your changes (when working in non-Workflow mode) or submit and approve the activity (when working in Workflow mode). If you discard the activity, any locks generated by the activity are also discarded. For more information about workflow modes, see Workflow and Activities Overview, page 1-20.

Keep in mind that:

- Locks are based on the device name, not the IP address of the device. Therefore, we recommend that you avoid defining two devices with different names but the same IP address in Security Manager. Any attempt to deploy to both devices, especially at the same time, leads to unpredictable results.
- In addition, locks do not extend across different operations. For example, locking does not prevent one user from deploying to the same device that is being discovered by a different user.

Additional details about locking can be found in the following sections:

- Understanding Locking and Policies, page 5-9
- Understanding Locking and VPN Topologies, page 5-10
- Understanding Locking and Objects, page 5-10

Understanding Locking and Policies

Table 5-1 on page 5-9 summarizes the effects of policy locks in Security Manager.

Note The ability to modify policies and policy assignments is dependent on the user permissions assigned to the user. See the *Installation Guide for Cisco Security Manager*.

Table 5-1 Locking Summary

If Another User or Activity...	You Cannot...	You Can...
Changes a policy definition	• Modify the policy or assign it to other devices. • Unassign the policy (if it is a local policy)	Unassign the policy from any device (if it is shared).
Changes the definition of a rule-based policy with descendants	• Modify the parent policy or any of the descendants. • Assign the parent policy or any of its descendants to additional devices. • Change the rule inheritance of the parent policy or any of the descendants.	Unassign the policy from any device.
Changes a policy assignment without changing its definition	Modify the policy. **Note** In Policy view, a content lock is placed on the policy. In Device view, an assignment lock is placed on those devices whose assignment is being changed by the other user.	Assign and unassign the policy from other devices.
Changes a policy definition and changes its assignment	Modify the policy or assign it to other devices.	Unassign the policy from any device.

Related Topics

- Understanding Policy Locking, page 5-8
- Understanding Policies, page 5-1

Understanding Locking and VPN Topologies

If you change the device assignment for a VPN topology, or make changes to a specific VPN policy, a lock is placed on the whole VPN topology, and on any other topologies in which the policy is shared. This means that other users cannot make changes to the device assignment, nor can they make changes to any of the VPN policies defined for those VPN topologies.

In order to view and modify site-to-site VPN policies, you must have the required permissions for each device in the VPN topology. You also need permissions to add a device to a VPN topology. If you have different levels of permissions to the devices in the VPN topology, the lowest permission level is applied to the entire topology. For example, if you have read/write permissions to the spokes in a hub-and-spoke topology, but read-only permissions to the device serving as the hub, you are granted read-only permission to the policies and devices in the hub-and-spoke topology. For more information about permissions, see *Installation Guide for Cisco Security Manager*.

Note Unassigning devices from a VPN topology also creates device locks in the VPN topology, which means that these devices cannot be deleted from the inventory. Other users cannot edit the device assignments for the topology until you deploy configurations to all affected devices, including those you remove. The device is not actually removed from the topology until you deploy configurations.

Related Topics

- Understanding Policy Locking, page 5-8
- Chapter 25, "Managing Site-to-Site VPNs: The Basics"

Understanding Locking and Objects

When you create or modify a reusable object, that object is locked to prevent other users from modifying or deleting the same object. Additional rules for object locking include:

- An object lock does not prevent you from modifying the definition or assignment of a policy that uses that object.
- The lock placed on a policy does not prevent you from making changes to an object that is included in the policy definition.
- You can change the definition of any object even if it is part of a policy assigned to a device to which you do *not* have permissions.
- When an object makes use of other objects (such as network/host objects and AAA server group objects), the lock on the object does not prevent another user from modifying those other objects. For example, when you modify a AAA server group object, the lock on that object does not prevent another user from modifying any of the AAA servers that make up the AAA server group.

When an object is locked, users who try to modify that object see a read-only version of the relevant dialog box. When you are working in Workflow mode, a message indicates which activity has locked the object.

Related Topics

- Understanding Policy Locking, page 5-8
- Chapter 6, "Managing Policy Objects"

Customizing Policy Management for Routers and Firewall Devices

When you manage Cisco IOS routers or ASA, PIX, or FWSM firewall devices, you have the option of selecting which policy types to manage with Security Manager and which policy types to leave unmanaged. Managing a policy type means that Security Manager controls the configuration of the policy and considers the information that it stores in its database about that policy to be the desired configuration. Security Manager does not configure unmanaged policy types, nor does it track configurations of these types that were configured using other methods. For example, if you decide not to manage SNMP policies, any SNMP configurations that you configured using CLI commands are unknown to Security Manager.

Caution If you use AUS or CNS to deploy configurations to ASA or PIX devices, be aware that the device downloads a full configuration from AUS or CNS. Thus, reducing the policies managed by Security Manager actually removes the configurations from the device. If you intend to deselect some ASA/PIX policies for management to use other applications along with Security Manager to configure devices, do not use AUS or CNS.

The ability to customize policy management on routers and firewalls makes it possible, for example, to use Security Manager to manage DHCP and NAT policies while leaving routing protocol policies, such as EIGRP and RIP, unmanaged. These settings, which can be modified only by a user with administrative permissions, affect all Security Manager users.

Unmanaged policies are removed from both Device view and Policy view. Any existing policies of that type, local or shared, are removed from the Security Manager database.

To customize policy management for routers and firewalls, select **Tools > Security Manager Administration > Policy Management** to open the Policy Management Page, page 11-64. The policy types are organized in folders, with router and firewall (which includes all ASA, PIX, and FWSM devices) handled separately. Select or deselect policy types as desired and click **Save**. Subsequent processing depends on whether you are changing a policy type to be managed or unmanaged:

- **Unmanaging a policy type**—If you unmanage a policy type, and any device of that type has that policy configured, you must unassign the policies before unmanaging them. Security Manager displays a list of all devices that have assigned policies of that type, including the policy name, device name, and the user or activity that has a lock on the policy. If you click **Yes** to continue unmanaging the policy, Security Manager obtains the required locks, unassigns the policies, and then unmanages the policy type.

  If a lock could not be obtained for even one device, no policies are unassigned, the policy type is not unmanaged, and you are told of the problem. You can then either manually unassign the policies from the affected devices, or release the user or activity locks, and try again to unmanage the policy type.

  Note Unmanaging a policy has no effect on the active configuration running on the device; Security Manager does not remove the configuration from the device. Instead, unmanaging the policy removes it from the database, and Security Manager no longer considers that part of the device configuration.

- **Managing a previously-unmanaged policy type**—If you start managing a policy type that you previously did not manage using Security Manager, it is possible that the active configuration on the device has commands controlled by the newly-managed policy type. **It is therefore important that you rediscover policies on all devices of that type (either all routers or all ASA, PIX, FWSM devices)**. This ensures that Security Manager has the current configuration for these policies.

If you do not rediscover policies and leave the newly-managed policies unconfigured, on the next deployment to the device, the existing settings configured on the device are removed. For more information on discovering policies on devices already managed, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Note Features that are unmanaged by Security Manager can still be modified manually with CLI commands or FlexConfigs. For more information about FlexConfigs, see Chapter 7, "Managing FlexConfigs".

Discovering Policies

Policy discovery enables you to bring your existing network configuration into Security Manager to be managed. Policy discovery can be performed by importing the configuration of a live device or by importing a configuration file. If you import a configuration file, the file must have been generated by the device (for example, by using the **show run** command on Cisco IOS Software devices); you cannot discover configuration files in any other format.

You can initiate policy discovery when you add a device by selecting the relevant options in the New Device wizard. For more information, see Adding Devices to the Device Inventory, page 3-7.

You can also initiate policy discovery on existing devices from Device view. For more information, see Discovering Policies on Devices Already in Security Manager, page 5-15.

When you initiate policy discovery on a device, the system analyzes the configuration on the device and then translates this configuration into Security Manager policies and policy objects so that the device can be managed. Warnings are displayed if the imported configuration completes only a partial policy definition. If additional settings are required, you must go to the relevant page in the Security Manager interface to complete the policy definition. Warnings and errors are also displayed if the imported configuration is invalid.

After performing policy discovery, you must submit your changes (or approve your activity when working in Workflow mode) to have the information included in change reports and to make the information available to other users. If you make any changes to the discovered policies, you must deploy the changes to the device for them to take effect. For more information, see Chapter 8, "Managing Deployment".

Tip Use the Security Manager Administration window to configure discovery-related settings that apply to all devices. For more information, see Discovery Page, page 11-25.

Policy Discovery and VPNs

In addition to performing discovery on individual devices, Security Manager allows you to discover the VPNs that are already deployed in your network. How you discover VPNs depends on the type of VPN being discovered:

- Site-to-Site VPNs—A wizard walks you through the discovery procedure step by step. For more information, see Site-To-Site VPN Discovery, page 25-20.

Tip We recommend that you deploy to a file immediately after discovering a Site-to-Site VPN. This enables Security Manager to assume full management of the relevant CLI commands that are configured on the device.

- IPSec and SSL Remote Access VPNs—You can discover IPSec and SSL VPNs when you discover policies on the device, either when you add the device to the inventory or if you discover policies on a device already in the inventory. Policies related to these VPNs are treated as regular device policies. However, when selecting discovery options, you must specifically select to discover RA VPN policies. For more information about remote access VPN policy discovery, see Discovering Remote Access VPN Policies, page 30-12. For more information about performing policy discovery, see Adding Devices to the Device Inventory, page 3-7 and Discovering Policies on Devices Already in Security Manager, page 5-15.

Note If you add a device using a configuration file, and discover security policies while adding the device, Security Manager cannot successfully discover policies that require that files be downloaded from the discovered device. This especially affects devices that include the **svc image** command in an SSL VPN configuration. Because Security Manager does not have the referenced file in its database, the **no** form of the command is generated for the discovered configuration.

Policy Discovery and Cisco IOS Routers and Catalyst Devices

Security Manager supports a subset of the complete list of commands available in the Cisco IOS software, mostly centered on security-related commands. You can discover all supported Cisco IOS commands. Commands that are not supported are left in place unless they conflict directly with a policy configured in Security Manager. For more information about performing policy discovery on Cisco IOS routers, see Discovering Router Policies, page 61-3. For more information about performing policy discovery on Catalyst devices, see Discovering Policies on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-1.

Tip We recommend that you deploy to a file immediately after discovering a Cisco IOS router or Catalyst device. This enables Security Manager to assume full management of the relevant CLI commands that are configured on the device.

Policy Discovery and Firewall Security Contents

When you add a device that has security contexts, you should discover all contexts and policies at the same time; otherwise, you will have to discover policies for each context separately. When you add the device, select **MULTI** for Context and do not select Security Context of Unmanaged Device. (If you select this option, only the admin context is imported, and it has no relationship to other security contexts on the device; select this option only if you want to manage the security context independently from the parent device.) Depending on how you add the device, you might need to select the option to discover security contexts. During discovery, Security Manager identifies each security context and adds it as a separate device to the device list, appending the security context name to the end of the parent's name; for example, if the parent is pix_141, the admin context would be pix_141_admin. (You can control the naming convention for security contexts; for more information, see Discovery Page, page 11-25). You can create new security contexts, or delete existing contexts, as well as create and delete policies for those contexts.

If you create multiple security contexts on FWSM, which are contained in Catalyst 6500 devices, and you are running IOS software on the chassis, add the chassis device using the SSH credentials for the chassis. Then Security Manager can identify each FWSM on the chassis, and give you the option to add each of them. During FWSM discovery, Security Manager discovers the security contexts for each FWSM, including the policies for the FWSM and for each context. However, if you are using the Catalyst OS on the device, you must discover each FWSM individually.

For more information about adding devices to the inventory, see Adding Devices to the Device Inventory, page 3-7.

Policy Discovery and IPS Devices

When you discover policies on an IPS device, the virtual sensors defined on the device are also discovered along with the policies defined for the virtual sensors. If more than one virtual sensor uses the same policy, that policy is created as a shared policy and is assigned to the virtual sensors. Policies defined for a single virtual sensor, or only for the parent device, are created as local policies. You cannot discover policies just for an individual virtual sensor; you can discover policies only on the parent device. If policies are discovered on the parent device that are not assigned to any virtual sensors, those policies are created as shared policies that are not assigned to any device or virtual sensor.

After discovering an IPS device that contains virtual sensors, you must submit your changes to the database for the virtual sensors to appear in the device selector.

Policy Discovery and Object Groups

When you perform policy discovery, any object groups already configured on PIX, ASA, FWSM, and IOS 12.4(20)T+ devices are brought into Security Manager as policy objects. For more information about how Security Manager policy objects are translated into object groups and vice-versa, see How Policy Objects are Provisioned as Object Groups, page 6-106.

In addition, **object network** and **object service** configurations on ASA 8.3+ devices are brought into Security Manager as host, network, or address range network/host objects, or service objects (as opposed to service group objects). The only exception is that address range objects that have the same address for the start and end range are instead created as host network/host objects.

Note For IOS devices, any objects discovered that are used by access control lists that are discovered as ACL objects are subsequently replaced during deployment by the contents of the object. Object groups used with ACL objects are not preserved, although they are discovered as Security Manager policy objects.

Policy Discovery and Security Manager Policy Objects

When you perform policy discovery, Security Manager tries to reuse the policy objects that you have already created in Security Manager. Based on the contents of the device configuration, the following are the possible actions:

- Named policy objects in the configuration—Existing policy objects are reused if their content matches the configuration on the device.

  If the contents of the named policy object does not match, the policy object is reused and a device-level override is created if **Allow Device Override for Discovered Policy Objects** is selected on the Discovery administration page. For more information, see these topics:

  - Understanding Policy Object Overrides for Individual Devices, page 6-18
  - Discovery Page, page 11-25

- Unnamed policy objects in the configuration—Existing policy objects are used if their content matches the configuration on the device. You can control this behavior by changing the value of the **Reuse Policy Objects for Inline Values** setting on the Discovery administration page.
- You can discover objects that have the same definition as existing objects, regardless of the setting you have defined for detecting redundant objects. For more information about this setting, see Policy Objects Page, page 11-66.

For more information on policy objects, see Chapter 6, “Managing Policy Objects”.

Policy Discovery and Access Control Lists

Certain policies in Security Manager support only standard or only extended ACLs, even if both types are supported by the CLI. In such cases, policy discovery works as follows:

- If the Security Manager policy supports only extended ACLs (for example, firewall service policies), any standard ACLs configured on the device for that policy are imported as extended ACLs.
- If the Security Manager policy supports only standard ACLs (for example, SNMP traps on IOS routers), any extended ACLs configured on the device for that policy are imported as standard ACLs.

During the discovery process, Security Manager will show any inactive ACLs that are imported as disabled. If you later deploy these disabled ACLs, they are removed from the device configuration.

Related Topics

- Frequently Asked Questions about Policy Discovery, page 5-27
- Viewing Policy Discovery Task Status, page 5-22
- Understanding Policy Object Overrides for Individual Devices, page 6-18

Discovering Policies on Devices Already in Security Manager

When you add a device to the inventory, you can discover policies at the same time that you add the device. However, you can skip policy discovery and do it later, or rediscover policies after adding the device.

You might initiate policy discovery on existing devices when:

- You make changes to device configurations using CLI commands, for example, device upgrade. In such a situation, you can rediscover existing policies on the device to make sure that the Security Manager database has the most current information. We recommended you to enter out-of-band changes in Security Manager rather than performing a rediscovery. However, beginning with version 4.13, in a single discovery action, all the policies are properly discovered (as applicable for upgrade of ASA 8.x to 9.x).
- You want to discover a subset of policies (for example, platform-specific settings) that was not discovered when you first added the device to Security Manager.
- You want to import the factory-default configuration of a firewall device. For more information, see Default Firewall Configurations, page 46-2.

Caution If you perform policy discovery on a device *after* configuring policies in Security Manager but before you deploy your changes, the discovered policies overwrite the undeployed changes. For example, if you select the option to discover platform-specific settings, the discovered configuration overwrites any platform-specific undeployed policies you configured in Security Manager. This is true even if the discovered configuration does not include the specific platform policy you configured. For example, discovering platform-specific settings overwrites any routing policies that you have configured for the device in Security Manager, even if the configuration you discover does not contain any routing information. Another result of rediscovery is that any shared policies that were configured on the device are replaced by the local policies that are discovered.

Caution Under certain conditions, Security Manager may fail to discover ASA interfaces in system context. Specifically, if a rediscovery/deployment is done on the system context of a multiple context ASA without checking (selecting) "inventory," then Security Manager may fail to discover the interfaces on

other security contexts. This can potentially result in Security Manager altering or altogether deleting interface configurations of other contexts in a subsequent deployment. To avoid this problem, simply be sure to select "inventory" when doing a rediscovery of the system context.

Before You Begin

Ensure that no one is configuring policies on the device or deploying configurations to the device. If you rediscover policies on a device while a deployment job is deploying configurations to the device, you might not be able to see the deployed changes after the rediscovery. Use the Deployment Manager to determine if there are active jobs that include the device before you rediscover policies (select **Manage > Deployments**). If you inadvertently rediscover policies during a deployment job, wait until the deployment job is completed and then discover policies again to ensure that Security Manager is synchronized with the device.

Related Topics

- Viewing Policy Discovery Task Status, page 5-22
- Discovering Policies, page 5-12
- Frequently Asked Questions about Policy Discovery, page 5-27
- Understanding Policies, page 5-1
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Managing Shared Policies in Policy View, page 5-50

Step 1 Decide whether you need to discover policies on a single device or if you want to discover policies on more than one device at a time. Policy discovery options vary based on how you start the discovery process.

- **Single device discovery**—If you need to discover policies related to any of the following, you can do it using only single-device discovery. (Note that single-device discovery is the type of discovery performed when you add a device to the inventory.)
  - Security context configurations for ASA, PIX, and FWSM devices running in multiple context mode.
  - Virtual sensor configurations for IPS devices.
  - Service module information for Catalyst devices.
  - Policy discovery from a configuration file.
  - Policy discovery from the factory default configuration.
- **Bulk rediscovery**—If you need to discover policies for more than one device, you can perform bulk rediscovery. However, bulk rediscovery can be performed only on live devices (that is, devices currently running and accessible in your network), and you cannot discover security context, virtual sensor, or Catalyst service module configurations. (You can discover service modules if you select them directly instead of selecting the device that contains them.)

Step 2 **If you want to perform single-device discovery**, do the following:

a. In device view or map view, ensure that only one device is selected, then right-click and select **Discover Policies on Device**. This opens the Create Discovery Task dialog box.

Tip: If the dialog box is called Bulk Rediscovery, you need to close the dialog box and try again. Ensure that only a single device is selected and reissue the command. You must use the right-click menu; it is the only way to perform single-device discovery.

b. Modify the discovery task name, if desired, and select the following discovery options. For detailed information, see Create Discovery Task and Bulk Rediscovery Dialog Boxes, page 5-18.

- **Discover From**—Whether you are discovering from a live device (which is active and accessible in the network), a configuration file (click **Browse** to select the file on the Security Manager server), or factory default configuration (for ASA, PIX, and FWSM devices running an OS version for which a factory default configuration exists). You can discover the default configuration only for devices that run in single-context mode or for individual security contexts.

Tip: We recommend that you use the Factory Default Configuration settings when you add PIX, ASA, and FWSM devices manually (as described in Adding Devices by Manual Definition, page 3-26). You should discover the default configuration for single-context mode devices and for each security context on a multiple-context mode device. For more information about factory-default policies, see Default Firewall Configurations, page 46-2.

- **Discover Policies for Security Contexts**—Select this option for firewall devices running in multiple-context mode if you want to discover policies for the security contexts defined on them.

c. Select the types of policies you want to discover. For more information about the difference between different types of policies, see Service Policies vs. Platform-Specific Policies, page 5-2.

- Detect ASA-CX/FirePOWER Module—Determines if a CX or FirePOWER module is installed; see Detecting ASA CX and FirePOWER Modules, page 72-21 for more information.
- Inventory—Discovers basic device information (such as hostname and domain name), interfaces, and security contexts on devices running in multiple-context mode. On Cisco IOS routers, this option also discovers all interface-related policies, such as DSL, PPP, and PVC policies.
- Platform Settings—Discovers platform-specific policies, such as routing policies.
- Firewall Services—Discovers firewall services policies, such as access rules and inspection rules, on all platforms.
- NAT Policies—Discovers network address translation (NAT) policies, such as address pools, static translation rules, and dynamic NAT/PAT. Discovery of NAT policies is supported on ASA, ASA-SM, PIX and FWSM devices.
- Routing Policies—Discovers routing policies for ASA devices.
- SSL Policy—Discovers SSL policy for ASA devices.
- RA VPN Policies—Discovers IPSec and SSL remote access VPN policies, such as IKE proposals and IPsec proposals.
- IPS—Discovers IPS policies, such as signatures and virtual sensors.

For more information, see Create Discovery Task and Bulk Rediscovery Dialog Boxes, page 5-18

d. Click **OK**. The discovery task is initiated and the Discovery Status dialog box opens so you can view the task status (see Discovery Status Dialog Box, page 5-23). You cannot perform other tasks in Security Manager while discovery is in progress.

Step 3 **If you want to perform bulk rediscovery**, do the following:

a. In device view, do one of the following:

- Select a device group, or multiple devices, then right-click and select **Discover Policies on Device**. Ensure that the Bulk Rediscovery dialog box opens.

Tip: If the dialog box is called Create Discovery Task, you need to close the dialog box and try again. Ensure that a device group or more than one device is selected and reissue the command.

- Select **Policy > Discover Policies on Device**. This opens the Device Selector dialog box. Select the devices you want to discover from the Available Devices list and click **>>** to move them to the Selected Devices list. Click **Next**.

Note If you use the right-click command, Security Manager assumes you have selected the desired devices. You can always click the **Back** button to go to the Device Selector screen and change the device list.

b. Modify the discovery task name, if desired, and select discovery options. For detailed information, see Create Discovery Task and Bulk Rediscovery Dialog Boxes, page 5-18.

The devices are organized in groups according to device type, with your device groups (if any) shown within each type:

- To change options for all devices of a given type, select the device type folder and modify the Discover Device Settings options. If the Discover drop-down list shows Multiple Values, then there are different discovery options selected for devices of that type. If you change the value, it changes for all devices. The check boxes for the policy types (explained above for single-device discovery) are available only if you select Policies and Inventory. Only options available for all devices in the selected group are shown, so you might need to select individual devices separately to select the most appropriate set of options.
- To change options for a single device, click the + icons next to folders to open them until you find the device, select the device, and then select the discovery options.

Note When the list of options is not expanded, all policies, namely the Platform Settings, Firewall Policies, NAT Policies and RA VPN, are discovered. However, when you expand the list of options, the discovery will be based on the options that you select from the available list.

c. Click **Finish**. The discovery task is initiated and the Discovery Status dialog box opens so you can view the task status (see Discovery Status Dialog Box, page 5-23). You cannot perform other tasks in Security Manager while discovery is in progress.

Create Discovery Task and Bulk Rediscovery Dialog Boxes

Use the Create Discovery Task dialog box to have Security Manager discover the policies for a device that is already in the device inventory. Use the Bulk Rediscovery dialog box to discover policies on more than one device at a time. Your options for policy discover differ based on which dialog box you use. For detailed information on the procedure, including how to get to each of these dialog boxes, see Discovering Policies on Devices Already in Security Manager, page 5-15.

You can also discover policies when you add the device to the inventory. For more information about adding devices, see Adding Devices to the Device Inventory, page 3-7.

Navigation Path

In Device view, select a device from the Device selector and do one of the following:

- Select **Policy > Discover Policies on Device** to perform bulk rediscovery.
- Right-click the device in the Device selector and select **Discover Policies on Device**. If a single device is selected, you get the Create Discovery Task dialog box. Otherwise, you are performing bulk rediscovery.

Tip You can also right click a device in Map view and select **Discover Policies on Device**.

Related Topics

- Discovering Policies, page 5-12
- Viewing Policy Discovery Task Status, page 5-22
- Selecting or Specifying a File or Directory in Security Manager, page 1-50
- Discovery Status Dialog Box, page 5-23

Field Reference

Table 5-2 Create Discovery Task Dialog Box

Element	Description
Discovery Task Name	The name assigned to the discovery task. Security Manager automatically generates a name for the task based on the current date and time, but you can modify this name as desired.
Selected Devices table (Bulk rediscovery only)	The devices you selected for rediscovery. The devices are organized in groups according to device type, with your device groups (if any) shown within each type: • To change options for all devices of a given type, select the device type folder and modify the Discover Device Settings options. If the Discover drop-down list shows Multiple Values, then there are different discovery options selected for devices of that type. If you change the value, it changes for all devices. The check boxes for the policy types (explained above for single-device discovery) are available only if you select Policies and Inventory. Only options available for all devices in the selected group are shown, so you might need to select individual devices separately to select the most appropriate set of options. • To change options for a single device, click the + icons next to folders to open them until you find the device, select the device, and then select the discovery options. **Tip**: To change which devices are selected for rediscovery, click **Back** to go to the Device Selector dialog box.

Table 5-2 Create Discovery Task Dialog Box (continued)

Element	Description
Discover From Config. File (Not available for bulk rediscovery)	The source of policy information to be discovered: • Live Device—Discover policies directly from the device. • Config File—Discover policies from a configuration file. Specify the location of the file in the **Config File** field. Click **Browse** to select the file on the Security Manager server. You can discover policies only from configuration files that were generated from the device (for example, with the **show run** command). For more information, see Adding Devices from Configuration Files, page 3-22. • Factory Default Configuration—Performs discovery on a firewall device using a file containing the factory-default settings for that device. Security Manager automatically chooses the appropriate file for the selected device (shown in the Config File edit box). This option is available only if Security Manager has a default configuration for the OS version running on an ASA, PIX, or FWSM device. You can discover the default configuration only for devices that run in single-context mode or for individual security contexts. For more information, see Default Firewall Configurations, page 46-2.
Discover Policies for Security Contexts (Not available for bulk rediscovery)	Whether to discover policies for each security context that is configured on a firewall device running in multiple-context mode. This field applies only to PIX, ASA, and FWSM devices. When deselected, Security Manager treats the entire device as having a single set of policies configured in single-context mode. For more information about security contexts, see Chapter 59, "Configuring Security Contexts on Firewall Devices".

Table 5-2 Create Discovery Task Dialog Box (continued)

Element	Description
Policies to Discover (for single-device discovery) Discover Device Settings (for bulk rediscovery)	The policy types to discover on the selected device. **Note** For bulk rediscovery, from the Discover drop-down menu, choose **Policies and Inventory** to enable the following options, **Inventory Only** to discover the inventory without discovering other policy types, or **Detect ASA-CX/FirePOWER Module** to determine if a CX or FirePOWER module is installed without discovering other policies. If the drop-down list has **Multiple Values** selected, this means that the devices in the selected group have different discovery options selected. If you change the selection, your change applies to all the devices in the group. The discovery options are: • Detect ASA-CX/FirePOWER Module—Determines if a CX module or FirePOWER module is installed; see Detecting ASA CX and FirePOWER Modules, page 72-21 for more information. • Inventory—Includes device information such as the hostname and domain name, interfaces, and security contexts (for firewall devices running in multiple-context mode). On Cisco IOS routers, this option also discovers all interface-related policies, such as DSL, PPP, and PVC policies. • Platform Settings—Includes all platform-specific policies that can be configured on the selected device. • Firewall Services—Includes all firewall service policies. For more information, see Chapter 12, "Introduction to Firewall Services". • NAT Policies—Includes all network address translation (NAT) policies that are configured on the selected device, such as address pools, static translation rules, and dynamic NAT/PAT. Discovery of NAT policies is supported on ASA, ASA-SM, PIX and FWSM devices. For more information, see Chapter 24, "Configuring Network Address Translation". • Routing Policies—Discovers routing policies for ASA devices. For more information, see Chapter 56, "Configuring Routing Policies on Firewall Devices". • SSL Policy—Discovers SSL policy for ASA devices. • RA VPN Policies—Includes all IPSec and SSL remote access VPN policies that are configured on the selected device. This option is disabled if the device does not support remote access VPN configuration. For more information, see Chapter 30, "Managing Remote Access VPNs: The Basics". • IPS Policies—Includes all IPS policies that are configured on the selected device. For more information, see Overview of IPS Configuration, page 36-5 or Overview of Cisco IOS IPS Configuration, page 45-4.

Table 5-2 Create Discovery Task Dialog Box (continued)

Element	Description
	Notes: • Routing Policies and SSL Policy options are applicable for Adaptive Security Appliance (ASA) devices only. • If you select Platform Settings as the policies to discover, the Routing Policies and SSL Policy which are sub options of Platform Settings, cannot be deselected. • To discover either Routing Policies or SSL Policy or both, you can deselect the Platform Settings option and then select either Routing Policies or SSL policy or both to discover only those policies. • For non-ASA devices, the Routing Policies and SSL Policy options may be listed but will always be unavailable for selection. • In bulk rediscovery, for transparent mode and system context you can select the Routing Policies option but no discovery will happen.

Viewing Policy Discovery Task Status

When you initiate policy discovery a discovery task is created. For each policy discovery initiation, only one task is created regardless of the number of devices being discovered.

You can view the status of the current policy discovery task in the Discovery Status dialog box, which opens automatically when the task is initiated. This dialog box provides updated status information about the discovery task, including summary information about the task and details about each device being discovered.

You can abort a discovery task, if required. When you perform policy discovery on a single device, aborting the task results in partial discovery. In such cases, we recommend deleting the information and starting again. When you perform policy discovery on multiple devices, any devices for which discovery was completed before you aborted the operation are fully discovered. Security Manager automatically discards the information for any partially discovered device.

The Discovery Status dialog box also displays the appropriate warning and error messages if any problems are encountered during the discovery process. For example, if the CLI commands in a configuration file do not define a complete Security Manager policy, a warning message is displayed that you must complete the policy definition in the relevant Security Manager policy page.

For more information, see Discovery Status Dialog Box, page 5-23.

To view information about previous discovery tasks, select **Manage > Policy Discovery Status** to open the Policy Discovery Status window. Select the discovery task in the top pane of the window, and the results of the task are displayed in the lower panes. For more information about using the Policy Discovery Status window, see Policy Discovery Status Page, page 5-25.

Related Topics

- Discovering Policies on Devices Already in Security Manager, page 5-15
- Frequently Asked Questions about Policy Discovery, page 5-27
- Discovering Policies, page 5-12

Discovery Status Dialog Box

Use the Discovery Status dialog box to view detailed information about the current policy discovery task. The dialog box includes general information about the status of the task, as well as detailed information about any warnings or errors generated by the device being discovered.

The Discovery Status dialog box opens automatically when you initiate a discovery task on existing devices and when you add devices from the network, from a configuration file, or from an export file. For more information about initiating a discovery task, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Related Topics

- Viewing Policy Discovery Task Status, page 5-22
- Discovering Policies, page 5-12
- Adding Devices from the Network, page 3-12
- Adding Devices from Configuration Files, page 3-22
- Adding Devices from an Inventory File, page 3-31

Field Reference

Table 5-3 Discovery Status Dialog Box

Element	Description
Progress bar	Indicates what percentage of the discovery task on the current device has been completed.
Status	The current state of the discovery task.
Devices to be discovered	The total number of devices being discovered during this task. The number includes service modules, security contexts, and virtual sensors.
Devices discovered successfully	The number of devices discovered without errors.
Devices discovered with errors	The number of devices that generated errors during discovery.

Table 5-3 Discovery Status Dialog Box (continued)

Element	Description
Discovery Details table	The devices that are being discovered. Select a device to see the messages generated during the discovery of that device in the message list below the summary list. Besides the device name, information in the table includes: • Severity—The overall severity level of the discovery task. For example, if the discovery task completed successfully, an Information icon is displayed. If the task failed, an Error icon is displayed. • State—The current state of the policy discovery task for the selected device: – Device Added—The device has been added to Security Manager, but policy discovery has not yet started. – Discovery Started—Policy discovery has started. – Reading and Parsing Device Config—The policy discovery task is interpreting the device configuration. – Importing Objects—The policy discovery task is importing objects from the configuration. – Importing Policies—The policy discovery task is importing policies from the configuration. – Discovery Complete—Policy discovery has been completed successfully. – Discovery Failed—Policy discovery failed due to errors. • Discovered From—The source of policy information. For example, when discovering from a configuration file, this field displays the name and path of the file.
Messages list	The messages generated during the discovery for the selected device. Select a message to see detailed information in the fields to the right of the list.
Description	Additional information about the message selected in the message list.
Action	The steps you should take to resolve the described problem.
Generate Report button	Click this button to create a discovery status report for this job. The report is a PDF file, saved to your client system, that includes a summary of the job. You can use this report for your own purposes or to aid in troubleshooting a problem with Cisco TAC. For more information, see Generating Deployment or Discovery Status Reports, page 10-30.

Table 5-3 Discovery Status Dialog Box (continued)

Element	Description
Abort button	Aborts the discovery task. If you abort the task when performing policy discovery on a single device, the result is partial discovery of that device. In such cases, we recommend deleting the information (for example, by discarding the activity) and starting again. If you abort the task when performing policy discovery on multiple devices, Security Manager automatically discards the information for any partially discovered device. Devices for which discovery was completed before you aborted the operation are fully discovered.

Policy Discovery Status Page

Use the Policy Discovery Status page to view the status of previous policy discovery and device addition tasks.

Navigation Path

Select **Manage > Policy Discovery Status**.

Related Topics

- Viewing Policy Discovery Task Status, page 5-22

Field Reference

Table 5-4 Policy Discovery Status Page

Element	Description
Task Table	
The upper portion of the window lists the previous policy discovery or device addition tasks. Select a task to view detailed information about it in the lower portion of the window. The columns in the table provide overall status information for the task. When adding devices that contain security contexts, the context discovery appears as a separate Policy Discovery task.	
Name	The name of the discovery or device addition task. This might be a system generated name or a name you specified when rediscovering device policies.
Type	The type of task, either Policy Discovery (when you rediscover device policies) or Add Device (when you add a device using the New Device wizard and elect to discover policies).
Start Time	The time the task started.
End Time	The time the task stopped.

Table 5-4 Policy Discovery Status Page (continued)

Element	Description
Status	The overall status of the task. One of the following: • Completed successfully—The task succeeded. • Completed with errors—The task was partially successful. This could occur if all policies were not discovered or if the device was added but no policies were discovered. • Completed with warnings—The task was successful but a minor problem occurred. • Failed—The task failed. No polices were discovered or no device was added because of errors or because you stopped discovery.
Generate Report button	Click this button to create a discovery status report for the selected job. The report is a PDF file, saved to your client system, that includes a summary of the job. You can use this report for your own purposes or to aid in troubleshooting a problem with Cisco TAC. For more information, see Generating Deployment or Discovery Status Reports, page 10-30.
Refresh button	Click this button to refresh the task list to update the information if there are tasks running in the background or if new tasks were created.
Delete button	Click this button to delete the selected task from the database. Deleting old tasks does not affect the related devices or discovered policies.
Discovery Details or Import Details Tables These tables display the devices included in the selected task. The name differs depending on the type of task (Discovery Details for Policy Discovery tasks, Import Details for Add Device tasks). Select a device to see the messages generated during the task for that device in the message list below the table.	
Device	The name of the device. If the name is followed by (deleted), the device is no longer in the Security Manager inventory.
Config File (Import Details only)	The location of the configuration file. This field is displayed only if you are importing from a configuration file.
Task Type (Import Details only)	One of the following: • Import only—Adding devices to Security Manager. • Import and Discover—Adding devices and discovering policies and inventory, or adding devices and discovering policies.
Severity	An icon for one of the following is displayed: • Error—The device addition or policy discovery failed. • Information—The device was added successfully or policy discovery was successful.

Table 5-4 Policy Discovery Status Page (continued)

Element	Description
State Details	These fields have the same meaning, although different names are used in the Discovery Details and Import Details tables. The fields describe the status of the task for the device: • Device Added—The device was successfully added to the inventory. • Device Add Failed—The device was not added to the inventory. • Discovery Completed—Discovery succeeded and the discovered policies are added to the Security Manager database. • Discovery Failed—No polices were discovered because errors occurred.
Discovered From (Discovery Details only)	One of the following: • Live Device—Security Manager contacted the device to obtain configuration and policy information. • File—Security Manager obtained the configuration and policy information from a configuration file.
Messages list	The messages generated during the task for the selected device. Select a message to see detailed information in the fields to the right of the list. The severity icons have these meanings: • Error—A problem was detected. • Warning—A minor problem occurred during discovery. • Information—An informational message about the selected device.
Description	Additional information about the message selected in the message list.
Action	The steps you should take to resolve the described problem.

Frequently Asked Questions about Policy Discovery

These questions and answers describe how policy discovery processes your device configurations into Security Manager policies.

Question: How does policy discovery work?

Answer: After you select the device whose policies, settings, and interfaces (inventory) you want to discover, Security Manager obtains the running configuration (from live devices) or the supplied configuration (when discovering from configuration files) and translates the CLI into Security Manager policies and objects. The imported configuration is added to the Configuration Archive as the initial configuration for the device. After discovery, you can review the discovered policies and objects and decide whether to commit them to the database. If you dislike them, you can discard them instead. Please note that commit and discard affect all discovered devices as a group and cannot be implemented on a per-device basis.

Question: When should I discover policies?

Answer: Typically, you should discover policies when you add devices to Security Manager. However, if you are creating devices in Security Manager (instead of importing live devices or configuration files), you must perform policy discovery after adding the device. You should also perform policy discovery in order to synchronize Security Manager with any out-of-band changes that have been made to the device, for example through the CLI.

Question: How can I determine the results of the discovery?

Answer: When you initiate a discovery task, a window opens that shows you the discovery status and results. You can also view a history of discovery task results on the Policy Discovery Status page (select **Manage > Policy Discovery Status**).

Question: Does Security Manager show which commands are not discovered, and what can I do about them?

Answer: In the discovery status window, go to the Message Summary section, then select **Commands Not Discovered**. Any undiscovered commands are listed in the Description field. You can either remove the command from the device and repeat the discovery process, or continue. If you continue, Security Manager will remove the unsupported command in the next deployment.

If Security Manager does not support a command found on a device, the discovery is generally not aborted; however, if the device has any access control entries (ACEs) that refer to unsupported object groups, the discovery is aborted. Other error messages, such as **User groups not supported**, might also provide details about undiscovered commands. Read the information in the Action box for suggestions.

Question: How are discovered policies reflected in the user interface?

Answer: Security Manager converts the device commands into policies. There is no difference in appearance between a policy discovered from a device configuration and one defined directly in Security Manager.

Question: I am using Auto Update Server for my PIX or ASA devices. How do I discover policies?

Answer: If a device has a static IP address, you can discover policies from the device. If it has a dynamic IP address, you must discover policies from the device's configuration file (offline).

Question: I am using Cisco Secure ACS to manage authentication and authorization to Security Manager. How does this affect policy discovery?

Answer: You must add all managed devices to Cisco Secure ACS before you can perform policy discovery and manage these devices in Security Manager. This includes security contexts on PIX/ASA/FWSM devices. For more information, see the *Installation Guide for Cisco Security Manager*.

Question: What should I do after discovering VPN or router platform policies?

Answer: Due to the way these features are discovered, Security Manager does not assume management of discovered VPN and router platform policies until after it deploys them. This means that if you discover a router, unassign one of its policies and deploy, no commands are removed from the router's configuration. We recommend, therefore, that you perform deployment to a file immediately after discovering VPN or router platform policies, *before* you make any changes to those policies. After this initial deployment, you can reconfigure these policies and deploy your changes as required.

Question: If I discover policies on a device and then deploy the policies from Security Manager without changing them, what is the difference between the original configuration on the device and the one that exists after the deployment?

Answer: Typically, there will be no differences between the new configuration and your original one, assuming you set up FlexConfigs for any unsupported CLI commands. However, in certain cases minor changes might occur in your ACL or object-group naming schemes. For more information, see How Policy Objects are Provisioned as Object Groups, page 6-106. In addition, any discovered objects that are not being used by a policy are removed from the configuration. There can also be instances where the new configuration is functionally equivalent to the old one but does not use the same commands.

Question: How does Security Manager handle my current CLI naming schemes for ACLs and object groups?

Answer: When you discover policies from a device, Security Manager tries to use the same names you have used. However, depending on your naming scheme, some minor differences might occur between what you defined on your device and the policies created through discovery. Additionally, there is a possibility that a naming conflict can occur between an existing ACL or object on the device and the name required for the new policy or object; in this case, Security Manager generates a different name so as not to misconfigure the device. For example, if the name of a discovered object conflicts with an object of the same type that already exists in Security Manager, a suffix is added to the name of the new object to make it unique or a device-level override is created.

Question: Are all configuration commands discovered and brought into Security Manager?

Answer: No. Security Manager does not discover all device configuration commands. Instead, it discovers security policies. For any configuration commands not discovered, use the FlexConfig feature to include the commands that Security Manager does not support.

Question: If I rediscover policies on a device already in Security Manager, what happens to the policies assigned to the device?

Answer: If you rediscover policies on a device that you are already managing with Security Manager, the newly discovered policies replace the ones assigned to the device. All policies within the selected policy domain (firewall services, platform settings, or both) are replaced, not just the ones that are different on the device compared to the ones in the Security Manager database. If you assigned shared policies to the device, the assignment is removed and the shared policy is left unchanged (so that other devices that use the shared policy are not affected). After policy discovery, all policies assigned to the device are specific to that device; none of them are shared with other devices. If you want to use shared policies with the device, you must redo the assignments after policy discovery.

In addition, any customizations done to local policies are also lost. For example, if you used sections to organize rules-based firewall policies, the sections are removed and the rediscovered policy is a flat list of entries.

Question: Does Security Manager use existing policies and objects during policy discovery?

Answer: During policy discovery, Security Manager uses existing policy objects (ones that you already defined in Security Manager) when creating policies for the device. However, Security Manager does not reuse existing policies; all policies created during discovery are local to the device being discovered. Thus, you might find it beneficial to define your policy objects (such as network objects) before adding devices to Security Manager.

Question: After adding a device and discovering policies, I cannot submit my changes to the database; instead I get warnings such as “Connection Policies Not Set.” What must I do to complete the device addition?

Answer: When you add a device and discover policies (particularly when you add devices from configuration files), Security Manager warns you if the resulting configuration is incomplete in ways that will prevent it from successfully managing the device. Connection policies, for example, are simply the device credentials (user names and passwords) required to log into the device, as well as other connection-related configuration settings (such as HTTP settings). Because these missing settings result in an invalid configuration or prevent Security Manager from contacting and managing the device later, you are prevented from submitting the changes to the database. Ensure that you have complete and valid configurations for these settings, then resubmit your changes to the database.

Question: Why does the AAA policy not show the AAA configuration that I discovered on the device?

Answer: The AAA policy contains the default configurations for authentication, authorization, and accounting. Other AAA commands that specify a particular list name are mapped to the policies that reference them. If the list name is not referenced by a policy, it is not discovered.

Question: Why are parts of the AAA method list definitions configured on my router not discovered?

Answer: Security Manager does not support certain keywords, such as if-needed. Method lists containing these keywords are discovered without the keyword. If the default AAA definitions on the device contain unsupported keywords, the entire command is not discovered.

Question: Can I discover AAA servers on devices running IOS software that were configured using the server-private command?

Answer: Yes, you can discover these servers. However, Security Manager converts them into standard AAA servers that can be used globally or in multiple AAA server groups. The server-private command is not supported.

Question: What do I need to know about discovery and device hostnames?

Answer: When you discover a device, the hostname policy is populated with the hostname discovered on the device. However, the hostname listed in Device Properties is not updated with this value. Ensure that the hostname defined in the device properties is the correct DNS name for the device. For more information, see Understanding Device Properties, page 3-6.

Managing Policies in Device View and the Site-to-Site VPN Manager

You can use Device view or the Site-to-Site VPN Manager to manage both local policies and shared policies, as described in the following sections:

- Policy Status Icons, page 5-30
- Performing Basic Policy Management, page 5-31
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

To access Device view, select **View > Device View** or click the **Device View** button on the toolbar. To access the Site-to-Site VPN Manager, select **Manage > Site-to-Site VPNs** or click the **Site-to-Site VPN Manager** button on the toolbar.

Related Topics

- Understanding the Device View, page 3-1
- Managing Shared Policies in Policy View, page 5-50
- Understanding Policies, page 5-1

Policy Status Icons

You can learn the status of any policy in Security Manager at a glance by viewing the icon displayed next to the policy name.

Table 5-5 Policy Status Icons

Icon	Status
	The policy is not configured. Upon deployment, any policy of this type already present on the device is effectively removed.

Table 5-5 Policy Status Icons (continued)

Icon	Status
	A local policy is configured. The definition of this policy affects only the device or VPN topology on which it is configured.
	A shared policy is configured. Any changes to the definition of this policy affect all of the devices or VPN topologies to which this policy is assigned.
	A policy bundle is configured. Any changes to the definition of this policy affect all of the devices or VPN topologies to which this policy is assigned, whether those policies are assigned using the same policy bundle, another policy bundle that includes the shared policy, or are assigned the shared policy directly and not through a policy bundle.

Related Topics

- Understanding Policies, page 5-1

Performing Basic Policy Management

The following topics describe the operations you can perform on local policies in Device view. Local policies are policies that are specific to the device or VPN topology on which they are configured. They are not shared by other network elements.

- Configuring Local Policies in Device View, page 5-31
- Copying Policies Between Devices, page 5-33
- Unassigning a Policy, page 5-36. (This topic also applies to the Site-to-Site VPN Manager.)

Related Topics

- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Managing Shared Policies in Policy View, page 5-50
- Understanding Policies, page 5-1

Configuring Local Policies in Device View

Use Device view to configure local platform and service policies on individual devices. Each policy defines a particular configuration or security task that the device can perform, such as NAT, OSPF routing or inspection rules. Local policies are unnamed and are particular to the individual device on which they have been defined. Any changes that you make to a local policy do not affect other devices that Security Manager is managing.

When you configure a policy, a lock is placed on that policy to prevent other users from making changes to the same policy at the same time. See Understanding Policy Locking, page 5-8.

You can modify any local policy assigned to a particular device provided you have permissions to modify policies and to access that device. For more information about permissions, see the *Installation Guide for Cisco Security Manager*.

After configuring a policy, you must deploy the changes to the device in order to make them active on that device. For more information, see Chapter 8, "Managing Deployment"

Related Topics

- Understanding the Device View, page 3-1
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Copying Policies Between Devices, page 5-33
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

Step 1 In Device view, select a device from the Device selector, then select a policy for that device from the Device Policies selector. The details of the policy appear in the work area.

Step 2 Modify the definition of the policy as required. Click the Help button to access information specific to the selected policy. For more information, see:

- Chapter 25, "Managing Site-to-Site VPNs: The Basics"
- Chapter 30, "Managing Remote Access VPNs: The Basics"
- Chapter 12, "Introduction to Firewall Services"
- Overview of IPS Configuration, page 36-5
- Overview of Cisco IOS IPS Configuration, page 45-4
- Chapter 61, "Managing Routers"
- Chapter 46, "Managing Firewall Devices"
- Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers"

Step 3 Click **Save** to save your changes.

If this is the first time you are configuring this policy on this particular device, the icon next to the selected policy changes to indicate that the policy is configured and assigned locally to the device. For more information about policy status icons, see Policy Status Icons, page 5-30.

After you save the policy, the policy is configured but you are the only one who can view the changes. There are additional steps to take to commit your changes and to deploy them to the device. The exact changes depend on whether you are working in Workflow or non-Workflow mode. Before taking the additional steps, configure all of the policies that you want to deploy; you are not required to deploy policy changes one at a time.

Following is a summary of the additional steps you need to take:

- Submit your changes. Submission updates the database on the Security Manager server with your changes.
  - In non-Workflow mode, you submit changes by selecting **File > Submit**. You can also submit your changes and deploy them in a single step by selecting **File > Submit and Deploy**.
  - In Workflow mode, if you are working with an activity approver, you submit your activity, and the changes are committed when the activity is approved. If you are not working with an activity approver, your changes are committed when you approve your own activity. For more information, see Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20 and Approving or Rejecting an Activity (Workflow Mode), page 4-21.

  In both Workflow and non-Workflow mode, policies are validated when you submit them. For more information on validation, see Validating an Activity/Ticket, page 4-18.
- Deploy your changes. Deployment either updates the devices directly with the new configuration, creates configuration files that you can deploy yourself, or copies the configuration files to an intermediate server (Auto Update Server, Configuration Engine, or Token Management Server) from which the device retrieves the updates. The method you use depends on the requirements of

your organization, and you can select different methods for each device. For general information about deployment, see Working with Deployment and the Configuration Archive, page 8-25. For the specific steps based on workflow mode, and information on the deployment methods, see the following topics:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Deploying Directly to a Device, page 8-9
- Deploying to a Device through an Intermediate Server, page 8-10
- Deploying to a File, page 8-11

Copying Policies Between Devices

You can streamline device configuration by copying multiple policies, or even a complete set of policies, from one device to other devices that support the selected policies. This makes it easy, for example, to quickly configure a new firewall device with the same policies configured on an existing firewall device.

When you copy policies between devices, those policies that are local on the source device are copied locally to the target device. Shared policies assigned to the source device are copied as shared policies to the target device as well.

Tips

- If your intention is to assign a single shared policy to additional devices, we recommend that you use the assignment feature, rather than copying the policies. For more information about sharing policies in Device view, see Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.
- To create a new device of the same type that shares the same configuration and properties (including the operating system version, credentials, and grouping attributes) as the source device, use the Clone Device feature. For more information, see Cloning a Device, page 3-56.

Related Topics

- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Configuring Local Policies in Device View, page 5-31
- Understanding the Device View, page 3-1
- Policy Status Icons, page 5-30
- Filtering Items in Selectors, page 1-45

Step 1 In Device view, do one of the following:

- Select **Policy > Copy Policies Between Devices**. The Copy Policies wizard starts at step 1, the Copy Policies from this Device page. Select the device that has the policies you want to copy and click **Next**.

- Right-click the device in the Device selector, then select **Copy Policies Between Devices**. The Copy Policies wizard selects the device as the source device and starts at step 2, the Select Policies to Copy page. You can change the source device by clicking **Back**.

Tip You can also right click a device in Map view and select **Copy Policies Between Devices**.

Step 2 Select the policies you want to copy on the Select Policies to Copy page. Initially, most policies from the source device (both local and shared) that can be copied are selected. You can change the selection, however, if you select a policy that depends on another policy, you must select the dependant policies. Security Manager will prompt you if your selections are not valid.

Consider the following when selecting policies:

- Selecting the check box for a policy group selects all of the policies in that group.
- When you copy policies between firewall devices (ASA, PIX, FWSM), copying the failover policy automatically copies the interface policy and vice-versa.
- It is usually not a good idea to copy interface policies, because these policies can have specific IP addresses. Other types of policies that you should carefully consider before copying them include NAT, routing, or the IPS policy on IOS devices.
- If you select the security contexts policy (for FWSM, PIX Firewall, or ASA devices), you must submit your changes after copying the devices for the contexts to appear in the device selector. In non-Workflow mode, select **File > Submit**. In Workflow mode, submit your activity.

Step 3 Use the policy object copy options to determine how policy objects are handled. These options are not mutually exclusive, and the combination you select has important implications on how the policies are defined on the target devices.

These are the possible combinations and their meanings:

- To ensure that the target devices have the same policy object settings as the source device, select both **Copy the Global Values of Policy Objects** and **Copy the Overridden Values of Policy Objects**.
- To ensure that if a policy object is used on the target device, its value is not overridden, select **neither** option. If a selected policy uses a policy object, and an equivalent policy on the target device uses the same policy object, the policy object's value defined on the target device is preserved. If the target device does not use the policy object, it is copied to the target using the policy object's global value (any overrides on the source device are ignored).
- To ensure that any policy objects on the target device use the policy object's global values, select **Copy the Global Values of Policy Objects** but deselect **Copy the Overridden Values of Policy Objects**. If the source device includes policies that use policy objects, only policies that use global values for the policy objects are copied. If the target device has an equivalent policy that uses local values for the policy object, the local values are replaced by the policy object's global values.
- To ensure that only policy objects with local values on the source device are copied to the target device, deselect **Copy the Global Values of Policy Objects** but select **Copy the Overridden Values of Policy Objects**. If the source device includes policies that use policy objects, only policies that override the policy object's global values are copied. The target devices get the source device's override value for the policy object.

The following table shows the possible outcomes when copying policy objects depending on which of the two options are selected:

Source Device	Target Device	User Option	Target device (copy result)
Global Definition	No reference	Any	Global Definition
Global Definition	Global Definition	Any	Global Definition
Global Definition	Device-level override	Neither option selected	Retains target device's override
		Copy the Global Values of Policy Objects only	Global Definition
		Copy the Overridden Values of Policy Objects only	Retains target device's override
		Both options selected	Global Definition
Device-level override	No reference	Neither option selected	Global Definition
		Copy the Global Values of Policy Objects only	Global Definition
		Copy the Overridden Values of Policy Objects only	Uses source device's override
		Both options selected	Uses source device's override
Device-level override	Global Definition	Neither option selected	Global Definition
		Copy the Global Values of Policy Objects only	Global Definition
		Copy the Overridden Values of Policy Objects only	Uses source device's override
		Both options selected	Uses source device's override
Device-level override	Device-level override	Neither option selected	Retains target device's override
		Copy the Global Values of Policy Objects only	Retains target device's override
		Copy the Overridden Values of Policy Objects only	Uses source device's override
		Both options selected	Uses source device's override

Click **Next**.

Step 4 Select the target devices to which you want to copy policies on the Copy Policies to these Devices page. Selecting the check box for a device group selects all of the devices in that group.

The device selector displays only those devices that support all of the policies you selected to copy. If you do not see all of the devices to which you want to copy policies, you can return to the policy selection page and deselect the more restrictive policies, and use the wizard a second time to copy the restrictive policies to the subset of devices that support them.

The device list is empty if no other device in the inventory can support all selected policies.

Tip After selecting devices, Click the **Preview** button to view a summary of the policies that will be copied. The summary shows the selected devices, the policies that will be copied to them, and any overrides that will be created, updated, or deleted due to the copied policies.

Step 5 Click **Finish**. You are asked to confirm that you want to copy policies.

The policies are copied to the target devices. If the copy operation fails for any target device, the copy is undone for successful devices, and you are shown a list of reasons why the copy failed for each problem device. Typically, copy failures are because someone else has a lock on a policy or device, or you do not have the required permissions to a device.

Unassigning a Policy

If you unassign a policy that has already been deployed to a device, in most cases the values that are defined for the policy are erased, effectively removing the policy from the device's planned configuration. When you perform deployment, the configuration for this feature that already exists on the device is removed.

The exact behavior depends on the type of policy that you unassign:

- Firewall service policies—If you unassign a policy, Security Manager erases the policy from the device.
- VPN policies:
  - Site-to-site VPN policies—You cannot unassign mandatory site-to-site VPN policies from the devices in the topology. If you unshare a mandatory policy, Security Manager assigns default values to the affected device. If you unassign an optional policy, Security Manager erases the configuration from the device. For more information, see Understanding Mandatory and Optional Policies for Site-to-Site VPNs, page 25-6.
  - IPSec remote access VPN policies—If you unassign a policy, Security Manager erases the policy from the device, even if it is a mandatory policy. In most cases, deployment fails if you do not create a new definition for the mandatory policy. In those cases where deployment does not fail, the device will fail to establish VPN tunnels.
  - SSL VPN policies—If you unassign a policy, Security Manager erases the policy from the device.
- Catalyst 6500/7600 or Catalyst switch policies—Interface and VLAN policies cannot be shared or unassigned. If you unassign a platform policy (such as IDSM settings or VLAN access lists) Security Manager removes the policy from the device.
- IPS policies—For all IPS device and service policies, a default policy is assigned to the device.
- PIX/ASA/FWSM policies—Policies that you cannot share with other devices cannot be unassigned from the device on which they are created. This includes interface, failover, security context, and resource policies. For other policy types (such as timeout policies), Security Manager makes a best effort to restore the system default configuration on the device.

- IOS router policies—Core connectivity policies, such as basic interface settings and accounts and credentials policies cannot be unassigned from the device on which they are created. If you unassign a device access policy that was used to define the password for configuring the device, you might prevent Security Manager from configuring that device in the future. For more information, see User Accounts and Device Credentials on Cisco IOS Routers, page 63-14.

  If you unassign a VTY or console policy, Security Manager restores a default configuration to ensure continued communication with the device. For all other policy types, if you unassign the policy, Security Manager erases the configuration from the device.

Related Topics

- Configuring Local Policies in Device View, page 5-31
- Copying Policies Between Devices, page 5-33
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30

Step 1 Do one of the following:

- (Device view) Select the device that has a policy you want to unassign.
- (Site-to-Site VPN Manager) Select the VPN topology that has a policy you want to unassign.

Step 2 Right-click the local policy and select **Unassign Policy**.

Note You can unshare a policy only if you have the Assign privilege mapped to your role. Cisco Security Manager displays error message for authorization.

You are asked to confirm that you want to unassign the current policy.

Working with Shared Policies in Device View or the Site-to-Site VPN Manager

Sharing policies makes it possible to configure multiple devices with common policies, which provides greater consistency in your policy definitions and streamlines your management efforts. Any changes to a shared policy affect all the devices and VPN topologies to which the policy is assigned. This makes it easy, for example, to update all of your Cisco IOS routers with new quality of service policies by updating the shared Quality of Service policy assigned to these devices.

When working in Device view or the Site-to-Site VPN Manager, you can take a local policy (such as a policy created during device discovery) and share it. You can then assign the shared policy to as many devices or VPN topologies as you want (provided they are not locked by another user; see Understanding Policy Locking, page 5-8), and you can change these assignments at any time. You can also take these shared policies that were created from the local policy and add them to a policy bundle. For more information on policy bundles, see Managing Policy Bundles, page 5-57.

Tip If you have a device that you are using as a template for the creation of other devices, you can quickly create a policy bundle that can be used for device configuration based on the template device. To do so, first make all policies on the device shared policies (see Sharing Multiple Policies of a Selected Device, page 5-42), then create a policy bundle from those shared policies.

In addition, you can take a shared policy that is assigned to a device or VPN topology and turn it into a local policy for that particular device or topology. This enables you to create a special configuration that affects only that device or topology. Other devices or topologies assigned the shared policy continue to use the shared policy as before.

As an alternative to sharing local policies, you can create new shared policies and manage them at the network level using Policy view. For more information, see Managing Shared Policies in Policy View, page 5-50. After creating the shared policy and assigning it to devices or VPN topologies in Policy view, you can return to Device view or the Site-to-Site VPN Manager and perform additional operations on the policy as described in the sections that follow. Note that all shared policies that you create in Device view or the Site-to-Site VPN Manager automatically appear as shared policies in Policy view.

Tip In Device view or the Site-to-Site VPN Manager, if you edit a shared policy, your changes are applied to all devices or VPN topologies that share the policy. Thus, you do not need to go to Policy view to edit shared policies. You are warned when you try to edit a shared policy that this will happen, to ensure you do not inadvertently make a change to more devices or topologies than what you intend. If you need to change the policy for just one device or topology, you can unshare the policy before editing it, as described in Unsharing a Policy, page 5-43.

The following topics describe how to share policies and the operations that can be performed on them in Device view or the Site-to-Site VPN Manager:

- Using the Policy Banner, page 5-38
- Policy Shortcut Menu Commands in Device View and the Site-to-Site VPN Manager, page 5-40
- Sharing a Local Policy, page 5-41
- Sharing Multiple Policies of a Selected Device, page 5-42
- Unsharing a Policy, page 5-43
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Adding Local Rules to a Shared Policy, page 5-45
- Inheriting or Uninheriting Rules, page 5-47
- Cloning (Copying) a Shared Policy, page 5-47
- Renaming a Shared Policy, page 5-48
- Modifying Shared Policy Definitions in Device View or the Site-to-Site VPN Manager, page 5-49
- Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49

Related Topics

- Importing Policies or Devices, page 10-13
- Understanding Policies, page 5-1
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30

Using the Policy Banner

When you view a device policy in Device view, or a site-to-site VPN policy in the Site-to-Site VPN Manager, there is a banner above the content of the policy in the work area. The banner provides information about whether the policy is local to the device or a shared policy. For shared policies, the banner also indicates the number of devices that use the policy. For policies that allow inheritance, the banner includes information about inheritance.

Messages might appear below the banner to indicate the following:

- That the policy is locked by another user. You cannot save changes to the policy until the other user submits (and approves) the changes, cancels an edit, or discards the changes.
- That the shared policy was imported. Imported policies might be re-imported at some point if the policy is managed on a different server. Any changes that you make are eliminated if the policy is imported again. Before editing the policy, ensure that you understand the protocols used in your organization for policy management and importation. You can control whether this message appears using an option on the Tools > Security Manager Administration > Policy Management page (see Policy Management Page, page 11-64).

You can use the links in the banner to create shared policies, assign a shared policy, and configure policy inheritance. The following illustration shows an example of a device policy banner.

Figure 5-3 Policy Banner Example

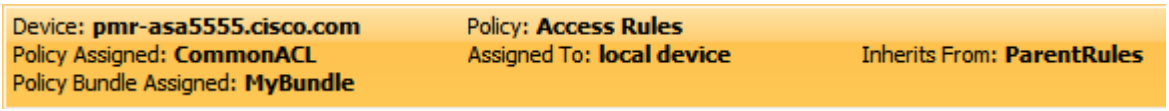

The fields in the policy banner have the following meanings and uses:

- Policy Assigned—The name of the policy assigned to this device or VPN. If the name has a link, you can assign a shared policy to the element by clicking the link. If there is no link, a shared policy cannot be assigned to this particular type of policy.
  - Local—The policy is a local policy (configured on this device only) rather than a shared policy.
  - Specific policy name—The shared policy is assigned to the device policy.
- Assigned To—If a shared policy is assigned, the number of devices or VPNs to which the policy is assigned. If no shared policy is assigned, **local device** or **this VPN** is indicated. If the name has a link, you can do the following:
  - Local Device or This VPN links—Click the link to create a shared policy from this local policy. You can then assign the shared policy to other devices or VPNs.
  - Number of Devices or VPNs links—Click the link to change the devices or VPNs assigned to the shared policy.
- Inherits From—The name of the policy from which this policy inherits rules. This field appears only for policies that allow inheritance. Click the link to specify a policy or set of policies from which the policy will inherit rules. For more information about inheritance, see Understanding Rule Inheritance, page 5-4.

  The field can contain these entries:
  - None—The policy does not inherit rules from any other policy.
  - Single policy name—The policy inherits rules from this policy.
  - Multiple policy names separated by > signs—The policy inherits rules from the indicated hierarchy of policies.
- Policy Bundle Assigned—The name of the policy bundle assigned to this device or VPN.

Related Topics

- Understanding Policies, page 5-1
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30

- Sharing a Local Policy, page 5-41
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Adding Local Rules to a Shared Policy, page 5-45
- Modifying Shared Policy Definitions in Device View or the Site-to-Site VPN Manager, page 5-49
- Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49
- Inheritance vs. Assignment, page 5-6
- Understanding Policy Locking, page 5-8
- Importing Policies or Devices, page 10-13

Policy Shortcut Menu Commands in Device View and the Site-to-Site VPN Manager

When you right-click a policy in Device view or the Site-to-Site VPN manager, you get a list of commands that you can use on the policy. The shortcut command list includes only those commands available for the selected policy, so the list differs according to your selection.

The available commands depend on whether the policy:

- Is unassigned.
- Contains a local policy for that specific device or VPN topology.
- Contains a shared policy that might be assigned to multiple devices or VPN topologies.
- Can be shared. There are no shortcut commands for policies that cannot be shared between devices or topologies.

The current status of each policy type is indicated by the icon displayed next to the policy name. See Policy Status Icons, page 5-30.

The following table provides a comprehensive list of the possible commands.

Table 5-6 Policy Shortcut Commands

Menu Command	Description
Commands available for both local and shared policies	
Assign Shared Policy	Assigns an existing shared policy to the selected device or VPN topology. If the policy is already assigned a shared policy, your selection assigns a new shared policy, replacing the existing selection. See Assigning a Shared Policy to a Device or VPN Topology, page 5-44.
Inherit Rules	Enables you to identify a shared policy from which to inherit rules, or to remove any inheritance from the child policy. Child policies inherit both the mandatory rules and default rules that are defined in the parent policy. See Inheriting or Uninheriting Rules, page 5-47.
Additional local policy commands	
Share Policy	Shares the local policy so that it can be assigned to other devices or VPN topologies. See Sharing a Local Policy, page 5-41.
Unassign Policy	Unassigns the policy from the device or VPN topology. When deployed, the configuration that corresponds to the settings defined in this policy is removed from the device or the devices in the topology. See Unassigning a Policy, page 5-36.

Table 5-6 Policy Shortcut Commands (continued)

Additional shared policy commands	
Unshare Policy	Creates a local copy of the shared policy and assigns it to the device or VPN topology in place of the shared policy. See Unsharing a Policy, page 5-43.
Edit Policy Assignments	Enables you to change which devices or VPN topologies are assigned to this policy, not just the device or VPN topology you are currently viewing. See Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.
Clone Policy	Creates a copy of a policy with a new name. Use this option to create a new policy with the same definition as the policy from which it was created, which you can then edit. See Cloning (Copying) a Shared Policy, page 5-47.
Rename Policy	Renames the selected policy. See Renaming a Shared Policy, page 5-48.

Sharing a Local Policy

As your network grows, you might decide to convert a local policy into a shared policy that you can assign to multiple devices or VPN topologies (see Local Policies vs. Shared Policies, page 5-3). Sharing a policy provides a streamlined management approach that ensures that all devices or topologies assigned to the policy are configured in a consistent manner. For example, if you configure a set of firewall inspection rules on a particular device, sharing that device's inspection rules policy makes it possible to assign that policy to other devices, eliminating the need to configure each device individually. See Assigning a Shared Policy to a Device or VPN Topology, page 5-44.

In addition, having a shared policy enables you to update the configurations of each assigned device or topology at one time, saving time and promoting greater consistency across your set of managed devices.

When you share a policy, you must name the policy. (Local policies do not have names, because they are associated with only a single device or topology.) This enables you to identify this policy when managing shared policies in Policy view.

Related Topics

- Understanding the Device View, page 3-1
- Policy Status Icons, page 5-30
- Using the Policy Banner, page 5-38
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Unsharing a Policy, page 5-43
- Adding Local Rules to a Shared Policy, page 5-45
- Sharing Multiple Policies of a Selected Device, page 5-42
- Inheriting or Uninheriting Rules, page 5-47
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

Step 1 In Device view or the Site-to-Site VPN Manager, select a policy from the Policies selector, then do one of the following:

- (Device view only) Select **Policy > Share Policy**.

- Right-click the policy and select **Share Policy**.
- Click the **local device/this VPN** link in the Assigned To field in the policy banner. A warning dialog box called Local Policies Cannot Be Assigned to Multiple Devices opens to inform you that you are viewing a local policy. Click **Share Policy** to continue.

The Share Policy dialog box is displayed.

Step 2 Enter a name for the shared policy and click **OK**.

Policy names can contain up to 255 characters, including spaces and special characters.

Sharing Multiple Policies of a Selected Device

With one procedure, you can share multiple policies configured on a particular device. When you perform this procedure, you can choose to share all the policies configured on the device or only some of them. For example, you can take all the firewall service policies defined on an ASA device and share them.

Initially, the resulting shared policies are assigned only to the device from which the procedure was performed. However, you can then assign these shared policies to additional devices as required. See Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.

This feature provides a convenient way to take the policies configured on a single device and use them as a template for configuring similar devices. For example, after you discover the devices at your branch offices, you can take all the local access rules that you have configured on a similar device and share them with a single procedure so that you can assign them to the branch office devices.

Tip You can use this procedure to make all policies on the device shared policies and then create a policy bundle from those shared policies. This policy bundle can then be used to quickly configure new devices based on the template device.

Tip To create a new device of the same type that shares the same configuration and properties (including device operating system version, credentials, and grouping attributes) as the source device, create a clone of the device. For more information, see Cloning a Device, page 3-56.

Related Topics

- Understanding the Device View, page 3-1
- Copying Policies Between Devices, page 5-33
- Sharing a Local Policy, page 5-41
- Unsharing a Policy, page 5-43
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Filtering Items in Selectors, page 1-45

Step 1 In Device view, do one of the following:

- Select **Policy > Share Device Policies**. The Share Policies wizard opens at the Share Policies from this Device page (step 1). Select the device whose policies you want to share and click **Next**.

- Right-click the device and select **Share Device Policies**. The Share Policies wizard opens at the Select Policies to Share page (step 2); you can click **Back** to go to step 1 and select a different device, if desired.

Tip You can also right click a device in Map view and select **Share Device Policies**.

Step 2 On the Select Policies to Share page, select all policies that you want to share. Initially, all shareable policies configured on the device, whether local or shared, are selected. Deselect the check box next to each policy that you do not want to share.

Following are some tips:

- Local policies that are not checked remain local to the selected device.
- If you select a policy that is already shared, Security Manager creates a copy of that policy using the name that you define in the wizard.
- Selecting the check box for a policy group selects all of the policies in that group.
- If a policy is configured on the device, but you cannot select it (the check box is solid grey), it is an unshareable policy.

Step 3 Enter a name for the shared policies. All policies are given the same name. You can later rename the individual policies. For more information, see Renaming a Shared Policy, page 5-48.

If you select a policy that is already shared, Security Manager creates a copy of that policy using this name.

Step 4 Click **Finish**. The selected policies become shared policies, which you can then assign to additional devices as needed. For more information, see Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.

Unsharing a Policy

When you unshare a shared policy assigned to a particular device or VPN topology, you create a copy that becomes a local policy for that device or topology. This means that any subsequent changes made to the local policy affect only this particular device or topology. Other devices or topologies assigned the original shared policy continue to use the shared policy as before.

Note You can unshare a policy only if you have the Assign privilege defined for your role. Cisco Security Manager displays error message for authorization.

Note You cannot unshare a policy that is assigned to a device as part of a policy bundle. You must either unassign the policy bundle from the device or remove the shared policy from the policy bundle that is assigned to the device.

For example, Security Manager might be managing a BGP routing policy called MyBGP, which is assigned to 20 routers. If you decide that one of the routers (Router1) requires a variation of this policy, you can select the device, unshare the policy, and make the changes you need for that router. From that point on, Router1 has a local BGP policy while the other 19 routers continue to use the original shared policy, MyBGP.

Related Topics

- Understanding the Device View, page 3-1
- Sharing a Local Policy, page 5-41
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Policy Status Icons, page 5-30

Step 1 In Device view or the Site-to-Site VPN Manager, select a policy from the Policies selector, then do one of the following:

- (Device view only) Select **Policy > Unshare Policy**.
- Right-click the selected shared policy, then select **Unshare Policy**.

Note You can unshare a policy only if you have the Assign privilege mapped to your role. Cisco Security Manager displays error message for authorization.

Step 2 Click **OK**. The shared policy is converted into a local policy for the selected device or VPN topology. The shared policy icon in the Policies selector is replaced by the local policy icon.

Assigning a Shared Policy to a Device or VPN Topology

You can replace any shareable policy (local or shared) assigned in Device view or the Site-to-Site VPN Manager with an existing shared policy of the same type. For example, if you have a local NAT policy assigned to a Cisco IOS router, you can assign a shared NAT policy in its place. Similarly, if a shared NAT policy was assigned to the router, you can replace it with a different shared NAT policy.

Tip You can use bundle shared policies together to make assigning those policies easier. For more information, see Managing Policy Bundles, page 5-57.

If you are assigning a shared policy to replace a local, rule-based policy (for example, an inspection rules policy), any local rules that you configured are replaced by the rules defined in the shared policy. A warning message gives you the opportunity to preserve the local rules by inheriting the rules of the shared policy instead of assigning the shared policy in place of the local policy. For more information, see Inheritance vs. Assignment, page 5-6.

Tip If you want to use the rules defined in the shared policy and still keep your local rules, we recommend that you select the Inherit Rules option instead of assigning the policy. For more information, see Inheriting or Uninheriting Rules, page 5-47.

Note You can also inherit IPS signature policies and signature event actions, but inheritance works differently than for rules-based policies. For more information, see Understanding Signature Inheritance, page 39-3.

Related Topics

- Understanding the Device View, page 3-1
- Using the Policy Banner, page 5-38
- Unassigning a Policy, page 5-36
- Adding Local Rules to a Shared Policy, page 5-45
- Copying Policies Between Devices, page 5-33
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

Step 1 In Device view or the Site-to-Site VPN Manager, select a policy from the Policies selector, then do one of the following:

- (Device view only) Select **Policy > Assign Shared Policy**.
- Right-click the policy in the Policies selector, then select **Assign Shared Policy**.
- Click the link in the Policy Assigned field in the policy banner.

The Assign Shared Policy dialog box is displayed if there are any shared policies available for assignment.

Step 2 Select a shared policy from the displayed list to assign to the device or VPN topology and click **OK**. If the policy does not allow inheritance, the shared policy is assigned to the selected device and you are finished.

Step 3 If the policy allows inheritance, you are warned that the shared policy will replace the current policy and given the option to inherit the rules instead with the Local Policy Will Be Replaced dialog box.

Your options are:

- Assign Policy—Assign the shared policy to replace the existing local policy. If you choose to assign, all local rules are removed and they cannot be retrieved.
- Inherit From Policy—Inherit the rules of the shared policy. If you choose to inherit, the inherited rules are added to the local rules that are already defined in the device's local policy. Use inheritance instead of assignment when the device needs to maintain the set of local rules already defined for it.

Tip You can select **Do not show this again** to save your selection and have it applied to all future times that you assign rule-based policies. Otherwise, you are prompted each time you assign policies so that you can make different selections based on the circumstances. If you select this option, you can turn it off by resetting it on the Customize Desktop administration settings page (see Customize Desktop Page, page 11-10).

Adding Local Rules to a Shared Policy

After you assign a shared rule-based policy, such as access rules, to a device, you can define additional rules in the policy that are local to that device. Selecting this option creates an inheritance relationship, where the policy defined on the device inherits rules from the shared policy while adding rules that affect only this particular device. For more information about inheritance, see Understanding Rule Inheritance, page 5-4.

Local rules that you add to a device do not affect the shared policy from which the device inherits its remaining rules. For example, if the shared policy Access_Rules_South is assigned to five devices and you define local rules on one of those devices, the access rules policy on that device consists of the rules defined in Access_Rules_South plus the local rules; the other four devices continue to use only the rules defined Access_Rules_South.

Before You Begin

Assign a shared, rule-based policy to the device as described in Assigning a Shared Policy to a Device or VPN Topology, page 5-44.

Related Topics

- Understanding the Device View, page 3-1
- Cloning (Copying) a Shared Policy, page 5-47
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Unsharing a Policy, page 5-43
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

Step 1 In Device view, select a device from the Device selector, then select a shared policy assigned to that device from the Device Policies selector. You must select a rule-based policy, such as access rules. The details of the policy appear in the work area.

Step 2 Do one of the following:

- Select **Policy > Add Local Rules**.
- Right-click the policy, then select **Add Local Rules**.

A message is displayed indicating that the policy on this device is now defined as a child policy that inherits rules from the shared policy. If the shared policy in turn inherits rules from a different shared policy, those rules are automatically inherited as well.

Note To change the parent policy from which this policy inherits rules, see Inheriting or Uninheriting Rules, page 5-47.

Step 3 Click **OK** to confirm. In the work area, headings are added for local mandatory and default rules in addition to the mandatory and default rules inherited from the shared policy.

In the Device Policies selector, the status icon changes to the icon for a local policy. For more information, see Policy Status Icons, page 5-30.

Step 4 Define local rules as required.

Tip If you assign a shared policy after adding local rules, both the inherited rules and your local rules are replaced with the selected shared policy.

Inheriting or Uninheriting Rules

This procedure describes how certain types of rule-based policies, such as access rules, can inherit rules from shared policies of the same type. Child policies inherit both the mandatory rules and default rules that are defined in the parent policy.

When working in Device view, you can then define additional rules that are local to the selected device. For more information, see Adding Local Rules to a Shared Policy, page 5-45.

You can edit rule inheritance from either Device view or Policy view.

Related Topics

- Understanding the Device View, page 3-1
- Managing Shared Policies in Policy View, page 5-50
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Understanding Rule Inheritance, page 5-4
- Inheritance vs. Assignment, page 5-6
- Understanding Policies, page 5-1
- Using the Policy Banner, page 5-38

Step 1 Select a local or shared rule-based policy in either Device view or Policy view, then do one of the following:

- Select **Policy > Inherit Rules**.
- Right-click the policy, then select **Inherit Rules**.
- (Device view only) Click the link in the Inherits From field in the policy banner.

The Inherit Rules dialog box is displayed, containing a list of all shared policies of the selected type, including any inheritance relationships among them.

Step 2 Select the policy from which to inherit rules, or select **No Inheritance** to remove any inheritance from the child policy. The name of the parent policy is displayed below the selector.

For example, if you select an access rules policy called West Coast, your access policy inherits the rules of the West Coast policy. If the West Coast policy is a child policy of another access rules policy called US, your policy inherits the properties of the West Coast policy, which in turn inherits the properties of the US policy.

Step 3 Click **OK** to save your definitions. The work area displays the inherited rules under the name of the parent policy and any local rules, if defined, under the name of the original shared policy.

Cloning (Copying) a Shared Policy

You can clone an existing shared policy. This provides a useful shortcut for creating a new policy that is similar to an existing one; after creating the clone, you can modify it as required.

If you clone a rule-based policy with inheritance, the new policy contains the same inheritance properties as the policy from which it was created. For more information, see Understanding Rule Inheritance, page 5-4.

Tip If you clone a policy in Device view or the Site-to-Site VPN Manager, the new policy is assigned to the selected device or VPN topology. If you want to clone a policy without changing policy assignments, make the clone in Policy view.

Related Topics

- Understanding the Device View, page 3-1
- Managing Shared Policies in Policy View, page 5-50
- Renaming a Shared Policy, page 5-48
- Deleting a Shared Policy, page 5-56

Step 1 Select a shared policy in Device view, Policy view, or the Site-to-Site VPN Manager, then do one of the following:

- (Device or Policy view only) Select **Policy > Clone Policy**.
- Right-click the shared policy, then select **Clone Policy**.

The Clone Policy dialog box is displayed.

Step 2 Enter a name for the new policy and click **OK**.

Names can contain up to 255 characters, including spaces and special characters.

Renaming a Shared Policy

You can rename a shared policy. The new name is immediately reflected in all devices and VPN topologies to which the policy is assigned.

Related Topics

- Understanding the Device View, page 3-1
- Managing Shared Policies in Policy View, page 5-50
- Cloning (Copying) a Shared Policy, page 5-47
- Deleting a Shared Policy, page 5-56

Step 1 Select a shared policy in Device view, Policy view, or the Site-to-Site VPN Manager, then do one of the following:

- (Device or Policy view) Select **Policy > Rename Policy**.
- Right-click the policy, then select **Rename Policy**.

The Rename Policy dialog box is displayed.

Step 2 Enter a new name for the selected policy and click **OK**.

Names can contain up to 255 characters, including spaces and special characters.

Modifying Shared Policy Definitions in Device View or the Site-to-Site VPN Manager

You can modify any shared policy in Device view or the Site-to-Site VPN Manager by selecting one of the devices or VPN topologies to which the policy is assigned, making the necessary changes, and then saving these changes to the Security Manager server. Any changes made to a shared policy in Device view or the Site-to-Site VPN Manager automatically affect all devices to which the shared policy is assigned.

Tip To apply your changes only to the device or VPN topology that you are modifying, you must first unshare the policy (see Unsharing a Policy, page 5-43). This action converts the policy to a local policy and prevents your changes from affecting other devices or topologies.

Related Topics

- Understanding the Device View, page 3-1
- Using the Policy Banner, page 5-38
- Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49
- Configuring Local Policies in Device View, page 5-31
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30

Step 1 Do one of the following:

- (Device view) Select the device that has a shared policy you want to modify.
- (Site-to-Site VPN Manager) Select the VPN topology that has a shared policy you want to modify.

Step 2 Redefine the policy as required.

Step 3 Click **Save**. You are asked to confirm that you want to save your changes, reminding you that the changes you made will be applied to all devices or topologies to which the policy is assigned.

Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager

You can modify the list of devices or VPN topologies assigned a particular shared policy as required. If you remove a device or topology from a policy assignment, that policy is effectively removed from the device's or topology's planned configuration. Upon deployment, any configuration of that type that exists on the device or topology is removed. For more information about the implications of unassigning a policy, see Unassigning a Policy, page 5-36.

Caution Use the policy assignment feature with care, as unassigning a policy removes that configuration from the device or topology and can have unintended consequences. For example, if you unassign a device access policy from a Cisco IOS router and then deploy that change, you might prevent Security Manager from configuring that device in the future (see User Accounts and Device Credentials on Cisco IOS Routers, page 63-14).

Policy assignment can also be modified from Policy view. For more information, see Modifying Policy Assignments in Policy View, page 5-54.

Related Topics

- Understanding the Device View, page 3-1
- Using the Policy Banner, page 5-38
- Assigning a Shared Policy to a Device or VPN Topology, page 5-44
- Unassigning a Policy, page 5-36
- Copying Policies Between Devices, page 5-33
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Inheriting or Uninheriting Rules, page 5-47
- Inheritance vs. Assignment, page 5-6

Step 1 In Device view or the Site-to-Site VPN Manager, select a shared policy from the Policies selector, then do one of the following:

- (Device view only) Select **Policy > Edit Policy Assignments**.
- Right-click the policy and select **Edit Policy Assignments**.
- Click the ***n*** device/VPN link in the Assigned To field in the policy banner.

Step 2 Modify the list of devices or VPN topologies to which the policy is assigned, as follows:

- To assign the selected policy to additional devices or topologies, select them from the Available Devices/VPNs list, then click **>>** to move them to the Assigned Devices list.
- To unassign the selected policy from devices or topologies, select them from the Assigned Devices/VPNs list, then click **<<** to return them to the Available Devices/VPNs list. Devices or topologies that are unassigned from the policy remove this policy from their running configuration during deployment.

Tip To assign a policy to all the devices in a device group, select the name of the device group, then click **>>**.

Step 3 Click **OK** to save your assignment changes.

Managing Shared Policies in Policy View

Use Policy view to globally manage all the shared policies configured in Security Manager. Unlike Device view, which you use to manage all the policies configured on a selected device, Policy view enables you to manage all shared policies of a particular type regardless of device.

Policy view enables you to:

- Create new shared policies.
- Edit any policy configuration.
- Modify the list of devices or VPNs to which shared policies are assigned.
- Delete shared policies that are not assigned to any devices or VPNs.

To access Policy view, select **View > Policy View** or click the **Policy View** button on the toolbar.

Figure 5-4 shows the main areas of Policy view.

Figure 5-4 Policy View

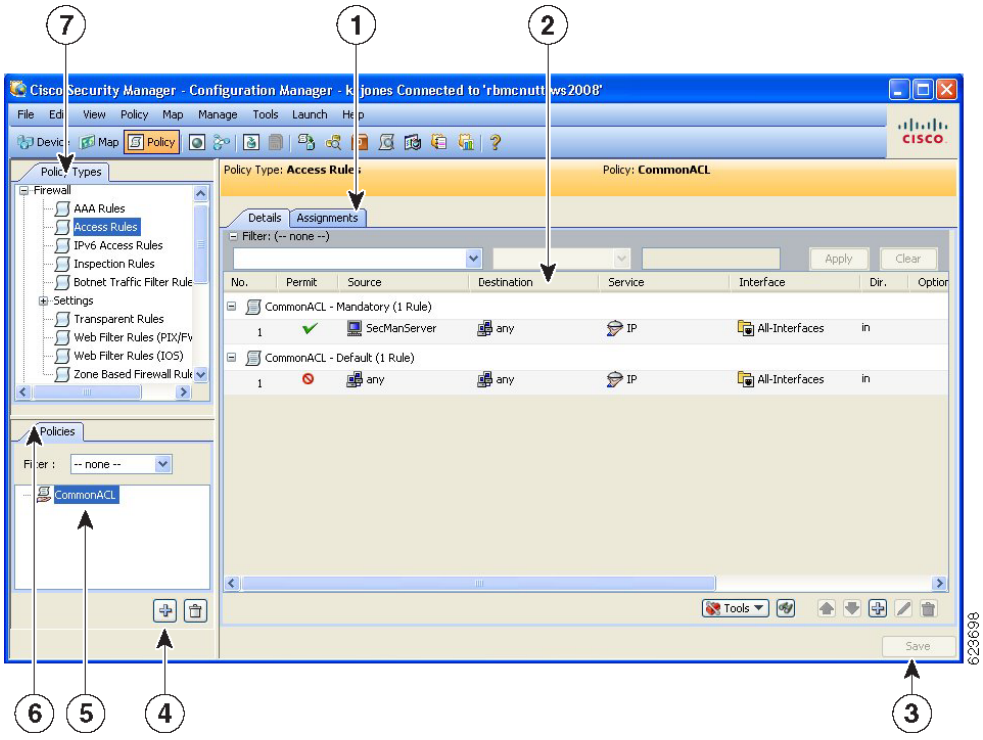

1	Assignments tab	5	Shared Policy selector
2	Work area and Details tab	6	Shared Policy filter
3	Save button	7	Policy Type selector
4	Create a Policy and Delete a Policy buttons		

- **(7) Policy Type Selector**—Lists the policy types available in Security Manager, divided by category. Clicking a policy type in the selector displays all the shared policies defined for that type in the Shared Policy selector. To create a new policy, right click the policy type and select **New [policy type] Policy** or click the Create a Policy button in the shared policy selector. For more information, see Policy View Selectors, page 5-52.
- **(4, 5, 6) Shared Policy Selector**—Lists the shared policies that are defined for the selected type. Clicking a policy in the selector displays the definition of the policy and its assignments in the work area. For more information, see Policy View Selectors, page 5-52.

  Right-click a policy in the selector to perform actions on the policy. For more information on the available commands, see Policy View—Shared Policy Selector Options, page 5-53.

  Use the Filter field to filter the list of policies displayed in the selector. For more information about creating filters, see Filtering Items in Selectors, page 1-45.
- **(1, 2, 3) Work Area**—Contains two tabs:

- Details—Use this tab to view and edit the definition of the selected policy. You can modify the definition as required; click **Save** in the work area to save your changes. Changes affect all devices or VPN topologies to which the policy is assigned. The information displayed on the Details tab is identical to the information displayed in Device view or the Site-to-Site VPN Manager and can be modified in exactly the same way. See Policy View Selectors, page 5-52.
- Assignments—Use this tab to view and edit the list of devices or VPNs to which a shared policy is assigned. For more information, see Modifying Policy Assignments in Policy View, page 5-54.

Related Topics

- Importing Policies or Devices, page 10-13
- Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37

Policy View Selectors

Policy view contains two selectors. The upper selector displays all the policy types available for a selected policy domain. The root of the policy type selector is the policy domain name. To display the policy types for a different policy domain, click the root of the tree and select a different domain from the list.

Policy domains include:

- Firewall—Lists all policy types for configuring firewall services. See Chapter 12, "Introduction to Firewall Services".
- NAT (PIX/ASA/FWSM)—Lists all NAT policies configured on PIX, ASA, and FWSM devices. See NAT Policies on Security Devices, page 24-15.
- NAT (Router)—Lists all NAT policies configured on Cisco IOS routers. See NAT Policies on Cisco IOS Routers, page 24-5.
- Site-to-Site VPN—Lists all policy types for configuring site-to-site VPNs. See Chapter 25, "Managing Site-to-Site VPNs: The Basics".
- Remote Access VPN—Lists all policy types for configuring remote-access IPSec and SSL VPNs. See Chapter 30, "Managing Remote Access VPNs: The Basics".
- Catalyst Platform—Lists all policy types for configuring Catalyst switches and 7600 routers. See Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers".
- IPS—Lists all policy types for configuring IPS devices. See Overview of IPS Configuration, page 36-5.
- IPS (Router)—Lists all policy types for configuring Cisco IOS IPS policies on IOS routers. See Overview of Cisco IOS IPS Configuration, page 45-4.
- PIX/ASA/FWSM Platform—Lists all policy types for configuring PIX/ASA/FWSM platform-specific policies. See Chapter 46, "Managing Firewall Devices".
- Router Interfaces—Lists all policy types for configuring platform-specific Cisco IOS router interface policies. See Chapter 61, "Managing Routers".
- Router Platform—Lists all policy types for configuring platform-specific Cisco IOS router policies. See Chapter 61, "Managing Routers".
- FlexConfigs—Lists all FlexConfig policies. See Chapter 7, "Managing FlexConfigs".

You can expand and collapse the selector as required to view all the available policy types and subtypes. To create a new policy, right click the policy type and select **New [policy type] Policy** or click the Create a Policy button in the shared policy selector.

Selecting a policy type from the Policy Type selector displays all the shared policies of that type in the Shared Policy selector. Local policies configured in Device view are not displayed.

For example, when you select a configuration policy type, such as NAT translation rules, the Shared Policy selector displays a flat list of each shared policy of that type. If you select a rule-based policy type, such as firewall access rules, the Shared Policy selector displays a hierarchical tree of shared policies. This enables you to view the inheritance relationships among the various policies. The Shared Policy selector includes a shortcut menu with options for actions that can be performed on that policy, such as renaming it.

Tip You can create and apply a filter to shorten the list of policies displayed in the Shared Policy selector. For more information about filters, see Filtering Items in Selectors, page 1-45.

Policy View—Shared Policy Selector Options

Right-click a policy in the Shared Policy selector of Policy view to display a shortcut menu for performing functions on the selected policy.

Related Topics

- Policy View Selectors, page 5-52
- Managing Shared Policies in Policy View, page 5-50

Field Reference

Table 5-7 Shared Policy Selector Options

Menu Command	Description
Clone Policy	Creates a new shared policy with the same definition as the policy from which it was cloned. See Cloning (Copying) a Shared Policy, page 5-47.
Rename Policy	Renames the selected policy. See Renaming a Shared Policy, page 5-48.
Add to Policy Bundle	Allows you to add the selected shared policy to a policy bundle. See Managing Policy Bundles, page 5-57.
Inherit Rules	Applies only to rule-based policies such as access rules. Causes a rule-based policy to inherit the rules of a different shared policy of the same type. See Inheriting or Uninheriting Rules, page 5-47.
New [policy type] Policy	Creates a new shared policy of the selected type. See Creating a New Shared Policy, page 5-54.
Delete Policy	Deletes the selected shared policy. See Deleting a Shared Policy, page 5-56

Creating a New Shared Policy

Use Policy view to create a new shared policy. In most cases, the new policy starts out undefined, but in certain cases (for example, many site-to-site VPN policies, such as IPsec proposals and GRE modes) default values are supplied. In all cases, the new policy is not initially assigned to any devices. If the new policy is a rule-based policy that supports inheritance, it can be created as a child of an existing shared policy of the same type. For more information, see Understanding Rule Inheritance, page 5-4.

Tip You can also create shared policies by converting local policies in Device view. For more information, see Sharing a Local Policy, page 5-41.

Related Topics

- Importing Policies or Devices, page 10-13
- Managing Shared Policies in Policy View, page 5-50
- Deleting a Shared Policy, page 5-56

Step 1 In Policy view, select a policy type in the Policy Type selector.

Step 2 Do one of the following:

- Right-click the policy type in the Policy Type selector, then select **New [policy type] Policy**.
- Right-click a policy in the Shared Policy selector, then select **New [policy type] Policy**.
- Click the **Create a Policy** button beneath the Shared Policy selector.

The Create a Policy dialog box is displayed.

Step 3 Enter a name for the new policy. Policy names can contain up to 255 characters, including spaces and special characters.

When creating a Translation Rules policy for NAT rules on security devices (PIX/ASA/FWSM), you must also choose a device software Version: **PIX/ASA 6.3-8.2** or **ASA 8.3 & Later**.

Step 4 Click **OK**. The new policy appears in the Shared Policy selector.

To configure a definition for the new shared policy, click the Help button in the toolbar with the Details tab open to see information specific to the type of policy you are creating. To assign the new shared policy, see Modifying Policy Assignments in Policy View, page 5-54.

Modifying Policy Assignments in Policy View

Use the Assignments tab in Policy view to modify the list of devices or VPN topologies to which you assigned a selected shared policy. The Assignments tab shows a list of all devices that are currently assigned the selected shared policy. It also shows devices that are assigned the policy through inheritance.

Assigning a policy to a device or VPN overwrites any policy of the same type (local or shared) that was previously assigned to the device in Security Manager. When deployed, the newly assigned policy overrides any policy of the same type that is already configured on the device, whether it was configured using Security Manager or using another method, such as the CLI.

When you unassign a shared policy from a device or VPN topology, Security Manager removes the policy from the planned configuration of that device or VPN topology. When the configuration defined by the policy is deployed, any configuration of the same type that is already configured on the device (including the devices in the VPN topology) is removed. For more information, see Unassigning a Policy, page 5-36.

Therefore, if your intention when performing unassign is to assign a different shared policy to a particular device or VPN topology, it is important to select the replacement policy and perform the assignment before performing deployment.

Tip Assigning a replacement policy is particularly important when you use a device access policy to configure the enable password or enable secret password on a Cisco IOS router. If you unassign this policy and fail to define a different password in its place before deployment, Security Manager might be unable to configure this device in the future. For more information, see User Accounts and Device Credentials on Cisco IOS Routers, page 63-14.

Alternatively, you can return to Device view and replace the shared policy assigned to the device with a different shared policy. For more information, see Assigning a Shared Policy to a Device or VPN Topology, page 5-44.

Note If you unassign a mandatory site-to-site VPN policy, such as an IKE proposal policy, Security Manager automatically replaces it with a default policy. If you unassign a mandatory remote access VPN policy, you must manually configure a new policy of that same type or deployment will fail.

Related Topics

- Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49
- Managing Shared Policies in Policy View, page 5-50

Step 1 In Policy view, select a policy type from the Policy Type selector, then select a policy from the Shared Policy selector. For more information about using these selectors, see Policy View Selectors, page 5-52.

Step 2 Click the **Assignments** tab in the work area.

The Assignments tab shows a list of all devices that are currently assigned the selected shared policy. It also shows devices that are assigned the policy through inheritance.

Step 3 Modify the list of devices or VPNs to which the policy is assigned, as follows:

- To assign the selected policy to additional devices or VPNs, select one or more items from the Available Devices/VPNs list, then click **>>** to move them to the Assigned Devices/VPNs list.

Tip To assign a policy to all the devices in a device group, select the name of the device group, then click **>>**.

- To unassign the selected policy from devices or VPNs, select one or more items from the Assigned Devices/VPNs list, then click **<<** to return them to the Available Devices/VPNs list.

Note Prior to the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, separate pages, policies and policy objects were provided for configuring IPv4 and IPv6 firewall rules and policies. With Security Manager 4.4 and ASA 9.0+, these policies and policy objects were combined or unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided in Device view, while in Policy view, IPv4 and unified versions of the AAA-, access- and inspection-rule policy types are provided.

If you assign an IPv4 AAA-, access-, or inspection-rule shared policy to a 9.0+ device, you will no longer be able to assign unified versions of those policies to that device. Likewise, if you assign a unified AAA-, access-, or inspection-rule shared policy to a 9.0+ device, you will no longer be able to assign IPv4 versions of those shared policies to that device--the device will not be included in the list of available devices on the Assignments tab for the shared policy.

Step 4 Click **Save** to save your assignment changes.

Deleting a Shared Policy

Use Policy view to delete a shared policy from Security Manager.

Before you delete a shared policy, you should unassign it from any devices that use it, and configure replacement policies for those devices. If a shared policy is assigned to a device, when the policy is deleted the device no longer has a policy configured for the deleted shared policy, other than whatever defaults might exist for the policy type. For more information about removing assignments, see Modifying Policy Assignments in Policy View, page 5-54.

Note If a shared policy is part of a policy bundle that is assigned to a device, you must remove the assignment before you can delete the shared policy.

Related Topics

- Creating a New Shared Policy, page 5-54
- Cloning (Copying) a Shared Policy, page 5-47
- Managing Shared Policies in Policy View, page 5-50

Step 1 In Policy view, select a policy type from the Policy Type selector, then select the policy to delete from the Shared Policy selector. For more information about using these selectors, see Policy View Selectors, page 5-52.

Step 2 Do one of the following:

- Right-click the policy, then select **Delete Policy**.
- Click the **Delete a Policy** button beneath the Shared Policy selector.

You are asked to confirm the deletion.

Managing Policy Bundles

Policy bundles are collections of shared policies that can be managed as a group. Policy bundles make managing shared policies easier by allowing you to create the bundle one time and then assign all of the policies in the bundle to a new device at once. The shared policies that are part of the bundle function in the same way as other shared policies and modifying any of the shared policies that are part of a bundle affects all devices that are assigned that policy either directly or through a policy bundle.

When creating a policy bundle, you can only assign one shared policy of each type to the policy bundle. Multiple policy bundles can be assigned to a device as long as the policy types in those policy bundles do not overlap.

When assigning a policy bundle to a device, if local policies on that device are the same policy type as those contained in the policy bundle, you are given the option to inherit or replace the existing policies.

Note When you unassign a policy bundle, all policies that are part of that bundle are removed from the device. Local policies will be lost and cannot be retrieved.

This section contains the following topics:

- Creating a New Policy Bundle, page 5-57
- Cloning a Policy Bundle, page 5-58
- Renaming a Policy Bundle, page 5-58
- Assigning Policy Bundles to Devices, page 5-59

Creating a New Policy Bundle

You can use the Policy Bundle view to create new policy bundles. When creating a policy bundle, you can only assign one shared policy of each type to the policy bundle.

Related Topics

- Managing Policy Bundles, page 5-57
- Cloning a Policy Bundle, page 5-58
- Renaming a Policy Bundle, page 5-58
- Assigning Policy Bundles to Devices, page 5-59

Step 1 Use one of the following methods to create a policy bundle:

- In Policy Bundle view, do one of the following:
  - From the All Shared Policies view, select the shared policies that you would like to bundle, then right-click on a selected shared policy and select **Create Policy Bundle**.
  - Right-click an existing policy bundle in the Policy Bundle selector, then select **Create Policy Bundle**.
  - Click the **Create a Policy Bundle** button beneath the Policy Bundle selector.
- To create a new policy bundle that includes all of the shared policies on a device, right-click a device in the Device selector in Device view, then select **Create Policy Bundle**.

The Create Policy Bundle dialog box is displayed.

Step 2 Enter a name for the new policy bundle.

Step 3 Click **OK**.

The policy bundle is added to the list of policy bundles in Policy Bundle view.

Step 4 To configure the definition for a policy bundle, do any of the following:

- In Policy Bundle view:
  - To add shared policies to the bundle, select **All Shared Policies** in the Policy Bundle selector and then drag and drop the required shared policies onto the policy bundle.
  - To remove shared policies from the bundle, select the bundle in the Policy Bundle selector. Select the shared policy you want to remove on the Details tab of the Policy Bundle View window, and then click **Delete**.
- In Policy view, right-click the shared policy you want to add to a policy bundle, select **Add to Policy Bundle**, and then select the bundle to which you want to add the shared policy.

Cloning a Policy Bundle

You can use Policy Bundle view to create a new policy bundle by cloning an existing bundle.

Related Topics

- Managing Policy Bundles, page 5-57
- Creating a New Policy Bundle, page 5-57
- Renaming a Policy Bundle, page 5-58
- Assigning Policy Bundles to Devices, page 5-59

Step 1 In Policy Bundle view, right-click an existing policy bundle in the Policy Bundle selector, then select **Clone Policy Bundle**.

The Clone Policy Bundle dialog box is displayed.

Step 2 Enter a name for the new policy bundle.

Step 3 Click **OK**.

The new policy bundle appears in the Policy Bundle selector.

Renaming a Policy Bundle

You can rename existing policy bundles from the Policy Bundle view. Renaming a policy bundle will not affect device assignments.

Related Topics

- Managing Policy Bundles, page 5-57
- Creating a New Policy Bundle, page 5-57
- Cloning a Policy Bundle, page 5-58

- Assigning Policy Bundles to Devices, page 5-59

Step 1 In Policy Bundle view, right-click an existing policy bundle in the Policy Bundle selector, then select **Rename Policy Bundle**.

The Rename Policy Bundle dialog box is displayed.

Step 2 Enter a new name for the policy bundle.

Step 3 Click **OK**.

The policy bundle name is updated in the Policy Bundle selector.

Assigning Policy Bundles to Devices

You can modify the list of devices assigned a particular policy bundle as required. Multiple policy bundles can be assigned to a device as long as the policy types in those policy bundles do not overlap. When assigning a policy bundle to a device, if local policies on that device are the same policy type as those contained in the policy bundle, you are given the option to inherit or replace the existing policies.

Note If any of the policies that are part of a policy bundle are not compatible with the device to which you are assigning it, the bundle cannot be assigned.

If you remove a device from a policy bundle assignment, all policies that are part of that bundle are effectively removed from the device's planned configuration. Local policies will be lost and cannot be retrieved. Upon deployment, any configuration of that type that exists on the device is removed. For more information about the implications of unassigning a policy, see Unassigning a Policy, page 5-36.

Caution Use the policy bundle assignment feature with care, as unassigning a policy bundle removes that configuration from the device and can have unintended consequences. For example, if you unassign a device access policy from a Cisco IOS router and then deploy that change, you might prevent Security Manager from configuring that device in the future (see User Accounts and Device Credentials on Cisco IOS Routers, page 63-14).

Related Topics

- Managing Policy Bundles, page 5-57
- Creating a New Policy Bundle, page 5-57
- Cloning a Policy Bundle, page 5-58
- Renaming a Policy Bundle, page 5-58

Step 1 In Policy Bundle view, select an existing policy bundle in the Policy Bundle selector.

The policy bundle details are displayed in the Policy Bundle main window.

Step 2 Click the **Assignments** tab.

Step 3 Modify the list of devices to which the policy bundle is assigned, as follows:

- To assign the selected policy bundle to additional devices, select them from the Available Devices list, then click >> to move them to the Assigned Devices list.

- To unassign the selected policy bundle from devices, select them from the Assigned Devices list, then click **<<** to return them to the Available Devices/VPNs list. Devices or topologies that are unassigned from the policy remove this policy from their running configuration during deployment.

Tip To assign a policy to all the devices in a device group, select the name of the device group, then click **>>**.

Step 4 Click **OK** to save your assignment changes.

The policy bundle name is updated in the Policy Bundle selector.

Managing Policy Objects

Policy objects enable you to define logical collections of elements. They are reusable, named components that can be used by other objects and policies. Objects aid policy definition by eliminating the need to define that component each time you define a policy. When used, an object becomes an integral component of the object or policy. This means that if you change the definition of an object, this change is reflected in all objects and policies that reference the object.

Objects facilitate network updates, because you can identify objects separately but maintain them in a central location. For example, you can identify the servers in your network as a network/host object called MyServers, and the protocols to allow on these servers in a service object. You can then create an access rule that permits the MyServers network/host object to send and receive traffic for the services defined in the service object. If a change is made to these servers, you need only update the network/host or service object and redeploy, instead of trying to locate and edit each rule in which the servers are used.

Objects are defined globally. This means that the definition of an object is the same for every object and policy that references it. However, many object types (for example, interface roles) can be overridden at the device level. Thus, you can create an object that works for most of your devices, yet customize the object to match the configuration of a particular device that has slightly different requirements. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18.

This chapter contains the following topics:

- Selecting Objects for Policies, page 6-2
- Policy Object Manager, page 6-4
- Working with Policy Objects—Basic Procedures, page 6-9
- Understanding AAA Server and Server Group Objects, page 6-27
- Creating Access Control List Objects, page 6-53
- Configuring Time Range Objects, page 6-71
- Understanding Interface Role Objects, page 6-73
- Understanding Map Objects, page 6-78
- Understanding Networks/Hosts Objects, page 6-80
- Understanding Pool Objects, page 6-92
- Configuring SAML Identity Provider, page 6-98
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- How Policy Objects are Provisioned as Object Groups, page 6-106

Selecting Objects for Policies

Modifying Policies using Drag and Drop

If you are modifying an existing policy, you can easily update the policy definition by dragging and dropping objects from the Policy Object Manager onto the applicable field in the policy. You can select a range of objects from the Policy Object Manager window by selecting the first object in the range and then, with the Shift key pressed, selecting the last object in the range. You can select multiple objects by clicking those objects while keeping the Ctrl key pressed. You can also select a range of objects and then add additional objects to your selection by using the Ctrl key method. To drag multiple objects, press and hold the Ctrl key while dragging or drag using the right-mouse button.

Creating Policies using Object Selector

When creating a policy, you often need to select one or more objects to include in the policy definition. For example, firewall policies make use of network/host objects, interface role objects, and service objects.

To include objects in policies, you can manually enter the object name or click the **Select** button to display an object selector dialog box. In certain cases, the object selector is prefiltered to display only the objects that are applicable to the policy that you are configuring. For example, when configuring a policy that requires a subnet, the object selector displays only those network/host objects that represent subnets, not network/host objects that represent single hosts. Object selectors make it easy for you to select which objects to include in a particular policy.

Additionally, object selectors enable you to create and edit objects of that type on the fly. This makes it easy to work with objects without leaving the policy you are defining to open the Policy Object Manager. For example, if when creating a dynamic NAT rule you discover that the ACL object you require does not exist, you can click the Create button to open the dialog box for creating an ACL object. When you finish creating the object, you are returned to the object selector with the new object selected and ready for inclusion in the policy. If you need to modify an existing object before using it, select it, click the Edit button and make your modifications, then click OK to save your changes; this returns you to the object selector.

When you create an object by opening the object editor from within a selector, the new object must conform to the requirements of the field from which the selector was opened. For example, if you open a selector from a field requiring a host and then decide to create a network/host object for that field, you must define the network/host object as a host.

There are two types of objects selectors—a simple list selector for policies that require you to select a single object, and a dual selector for policies that allow you to select multiple objects of a certain type. The following table explains these selectors and how to use them.

Table 6-1 Object Selectors

Element	Description
Type	The type of object to display in the selector, if there is an option. For example: • You can choose between network/host objects and interface roles when configuring sources and destinations in some rule-based policies. • You can choose between standard and extended ACL objects when configuring some ACLs (for example, when configuring VLAN ACLs on Catalyst 6500/7600 devices). **Tip** In some policies, if you select more than one type of object, they are displayed on different tabs within the field.
Available [object type]	Displays all objects that are relevant to the policy or object you are configuring. When selecting interfaces, be aware that there can be interfaces and interface roles with the same name. They can be distinguished by the icon displayed next to the name. For more information, see Specifying Interfaces During Policy Definition, page 6-76. **Tip** You can quickly find an object inside a selector by clicking in the list box and then starting to type the name of the object.
Selected [object type]	Displays the objects that you selected to apply to the policy or object that you are editing.
Multi-Object Selector Buttons	
>> button << button	Moves the selected objects from one list to the other list in the direction indicated. You can select multiple objects by using Ctrl+click. You can also move objects between lists by double-clicking them or by selecting them and pressing Enter.
Up/Down arrow buttons	For a limited number of object types, order matters. If the selector includes **Move Up** and **Move Down** buttons, arrange the objects in priority order. For example, when defining a method list for AAA, use the arrows to determine the order in which different types of AAA server groups are used.
Common Buttons	
Create button	Click this button to create an object of this type. **Tip** In a few cases, such as network/host and service objects, clicking this button opens a list from which you need to select a specific type for the object.
Edit button	Click this button to edit the selected user-defined object. If you try to edit a system-defined object, it is opened in read-only mode.

Related Topics

- Allowing a Policy Object to Be Overridden, page 6-18
- Filtering Items in Selectors, page 1-45

Policy Object Manager

Use the Policy Object Manager to:

- View all available objects grouped by object type.
- Create, copy, edit, and delete policy objects.
- Drag and drop objects onto existing policies to update the policy definition.
- Generate usage reports, which describe how selected objects are being used by other Security Manager objects and policies.

Navigation Path

In Device view or Policy view, click the **Policy Object Manager** button on the toolbar, or choose **Policy Objects** from the **Manage** menu. (The Policy Object Manager cannot be opened from Map view.)

Note When you open the Policy Object Manager, it is initially displayed as a pane in the lower half of the current view to make dragging and dropping objects easier. You can "undock" this pane, making the Policy Object Manager a separate window; you also can "re-dock" the window. See Policy Object Manager: Undocking and Docking, page 6-8 for more information.

Related Topics

- Creating Policy Objects, page 6-9
- Selecting Objects for Policies, page 6-2
- Generating Object Usage Reports, page 6-15
- Managing Object Overrides, page 6-17
- Filtering Tables, page 1-48

Field Reference

Table 6-2 Policy Object Manager Window

Element	Description
Object Type selector, or table of contents (Left pane)	Lists the object types available in Security Manager. When you select an object type, all existing objects of that type are listed in the table in the right pane. The objects are organized into three folders: Favorites, Recent Objects, and All Object Types. Click the arrow to left of a folder name to expand that folder. You can also specify your favorite object types and they will be presented in a separate list so that they can be more easily accessed. To add an object type to your favorites list, right-click the object and then select **Add to Favorites**. To remove an object type from your Favorites list, right-click the object and choose **Remove from Favorites**. Recent Objects is a list of the ten most recently modified objects. Click a recent object to see a summary of the object that includes the name, type, description, and last modified date. You can also access the View Object, Edit Object, and Find Usage buttons for the object. Expand the All Object Types folder to view all types of object available.

Policy Object Table (Right Pane)

The policy object table in the right pane lists existing objects of the type selected in the table of contents. Using this table, you create new objects and work with existing ones. You can use the buttons below the table, or right-click within the table to see additional commands (see Policy Object Manager Shortcut Menu, page 6-8).

Except for the Access Control Lists (ACL) object, there is one table per object type. For ACLs, there are tabs to separate Extended, Standard, Web, and Unified ACL objects. Select the appropriate tab to work with the desired object type.

The columns in the table vary based on the type of object you select. You can alter the columns displayed in the table by right-clicking the table heading and selecting or deselecting columns in the Show Columns command. You can also sort the information by the contents in a column by clicking the column heading; click the heading to toggle between alphabetical and reverse alphabetical sorting.

For detailed information on the settings that are displayed in the table, click the Create or Edit buttons below the table and click Help in the dialog box that is opened. The following section, “Table Columns,” is a description of the columns that you typically see.

Buttons Above Table	
Referenced	Select this option to view reference information for objects. When selected, a "Referenced" column is added to the table to display information on whether an object is being used by any policies or policy objects.
Find Usage	Use the Find Usage feature to view a report on the policies or policy objects that are using the selected object and any device overrides for the object. For more information, see Generating Object Usage Reports, page 6-15.
View Object	When a single object is selected in the table, you can click this button to open the Edit dialog box for that type of object in read-only mode to view the settings for that particular object.

Table 6-2 Policy Object Manager Window (continued)

Element	Description
Export	Use the Export feature to download a CSV file of the object data for the selected object type.
Print	Use the Print feature to print the object data for the selected object type.
Filter	Allows you to filter the rows displayed to help you find items in a large table. For more information, see Filtering Tables, page 1-48.
Table Columns	
*	Indicates the policy object status: • - Policy object has been locked for editing. Hover over the lock icon to see the user and ticket/activity that has the object locked. • - Policy object has been modified in the current ticket/activity but the changes have not been submitted. **Note** You can hover over the status icons to see details about the ticket/activity in which the policy object has been modified/locked and to navigate to that ticket/activity.
Icon (unlabeled field)	The icon displayed for a policy object type identifies objects of that type wherever they appear, such as in rules tables. If the icon includes the image of a pencil, you can edit it.
Name	The name of the policy object.
Content	A summary of the object definition that might not include all defined settings.
Permit	For ACL objects, if the Access Control Entry (ACE) allows traffic, a check mark appears in the Permit column. If the action is deny, a red circle with a slash appears.
Category	The category object that is assigned to the object, if any. Categories help you organize and identify rules and objects. For more information, see Using Category Objects, page 6-13.
Overrides	Whether a user can override the object properties at the device level. A check mark indicates that the object can be overridden. Not all object types are overridable. If an object has been overridden, the Overrides column displays the number of overrides for that object. You can click on the number to see the list of overrides. For more information about device overrides, see Managing Object Overrides, page 6-17.

Table 6-2 Policy Object Manager Window (continued)

Element	Description
Referenced	Whether the object is being used in any policy definitions. You can find out which policies or policy objects are using the selected object and any device overrides for the object using the Find Usage feature (see Generating Object Usage Reports, page 6-15). **Note** To view reference information, make sure the Referenced option is selected on the toolbar above the Policy Object Table. **Note** The Referenced column reports usage based on both committed data and uncommitted data across all activities/tickets, whereas as the Find Usage feature only reports usage based on committed data and data from the current activity/ticket.
Description	The description for the object. If the column is too narrow to display the description, you can double-click the icon to view the description or mouse-over the icon.
Last Ticket(s)	If ticketing is enabled, shows the Ticket ID of the ticket last used to modify the object. You can click on the If ticketing is enabled, shows the ticket(s) associated with last modification to the object. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Last Modified Date	Shows the date and time the object was last modified.
Buttons Below Table	
	Click the New Object button to create a new object. The same icon is used for any button that adds an item to a table. **Tip** In a few cases such as Networks/Hosts and Services objects, clicking this button opens a list from which you need to select a specific type for the object. Clicking this button opens a dialog box to create the object. Click the Help button in the dialog box for information on the selected object type. Also, see Creating Policy Objects, page 6-9.
	Click the Edit Object button to edit the selected object. The same icon is used for editing any object in a table. The dialog box used for editing the object is the same as the one used for creating the object. If you try to edit a system-defined default object, you are allowed only to view the object contents. Click the Help button in the dialog box for information on the settings. For more information, see Editing Objects, page 6-12.
	Click the Delete Object button to delete the selected object. You can delete only user-defined objects that are not currently being used in a policy or another policy object. For more information, see Deleting Objects, page 6-16.

Policy Object Manager: Undocking and Docking

Whenever you open the Policy Object Manager, it is initially displayed as a pane in the lower half of the current view to make dragging and dropping objects easier. You can "undock" this pane, making the Policy Object Manager a separate window, you can "re-dock" the window, and you can close the pane or window:

- To undock the Policy Object Manager from the current view in the Configuration Manager window, click the Undock Window button in the upper right corner of the pane's title bar.
- To put the floating window back as a pane in the Configuration Manager window, click the Dock Frame button in the upper right corner of the Policy Object Manager window.
- To close either the pane or the floating window, click the Close button in its upper right corner.

Navigation Path

In Device view or Policy view, click the **Policy Object Manager** button on the toolbar, or choose **Policy Objects** from the **Manage** menu. (The Policy Object Manager cannot be opened from Map view.)

Policy Object Manager Shortcut Menu

Right-clicking inside the policy object table in the Policy Object Manager, page 6-4 displays a shortcut menu for performing various functions on the selected object type.

Field Reference

Table 6-3 Policy Object Manager Shortcut Menu

Menu Command	Description
New Object	Choose this command to create a new policy object. Click Help in the dialog box that is opened for information specific to the object type. Also, see Creating Policy Objects, page 6-9. **Tip** For network/host and service objects, you need to also select an object type from the submenu.
Edit Object	Choose this command to edit the policy object selected in the table. If you select a system-defined default object, you are presented with a view-only look at the object definition. For more information, see Editing Objects, page 6-12.
Delete Object	Choose this command to delete the policy object selected in the table. You can delete only user-defined objects that are not being used in a policy or in another policy object. For more information, see Deleting Objects, page 6-16.
Enable/Disable Device Overrides	Choose the Enable Device Overrides command to enable device overrides on one or more devices on which overrides are disabled. Choose the Disable Device Overrides command to disable device overrides on one or more devices on which overrides are enabled.
Edit Device Overrides	Select this command to change the device-level overrides for this object using the Policy Object Overrides Window, page 6-20. You can create, edit, and delete overrides. For more information, see Managing Object Overrides, page 6-17.

Table 6-3 Policy Object Manager Shortcut Menu (continued)

Menu Command	Description
Clone Object	Select this command to create a copy of the policy object. For more information, see Cloning (Duplicating) Objects, page 6-14.
Copy Object	Choose this command to copy one or more selected objects to the system Clipboard. **Tip** You can also use **Ctrl+C** to copy objects.
Paste Object	Choose this command to paste the object(s) on the system Clipboard into another object. For example, you might add a host-type Networks/Hosts object to an existing group-type Networks/Hosts object. The two object types must be compatible. **Tip** You can also use **Ctrl+V** to paste objects.
Find Usage	Choose this command to generate a usage report for the selected object using the Object Usage dialog box. The usage report tells you where the object is currently being used. for more information, see Generating Object Usage Reports, page 6-15.
View Object	Choose this command to view the definition of the object using a read-only version of the edit dialog box for the object. For more information, see Viewing Object Details, page 6-14.

Working with Policy Objects—Basic Procedures

The following topics describe the actions that you can perform on policy objects. Some tasks are limited to certain types of objects. For example, not all types of object can be overridden, you cannot edit predefined objects, and you cannot import or export all objects.

This section contains the following topics:

- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Cloning (Duplicating) Objects, page 6-14
- Viewing Object Details, page 6-14
- Generating Object Usage Reports, page 6-15
- Deleting Objects, page 6-16
- Managing Object Overrides, page 6-17
- Importing and Exporting Policy Objects, page 6-23

Creating Policy Objects

Security Manager provides predefined policy objects of various types that you can use to define policies. Additionally, you can create your own objects, as required.

You can create objects in one of two ways:

- Using the Policy Object Manager window. This option is best suited for situations where you are defining one or more objects outside of the context of defining a particular policy. See Policy Object Manager, page 6-4.
- Using object selectors. When you define a policy that uses objects, object selectors include buttons for creating and editing objects so you don't have to leave the policy you are defining. This is frequently the best method to use, because during policy creation you are prompted for the specific type of object that applies to the situation, and you are more aware of the settings you need for the policy. See Selecting Objects for Policies, page 6-2.

Tip Your ability to create multiple objects with the same definition depends on a setting on the Policy Objects page in the Security Manager Administration window (select **Tools > Security Manager Administration**). By default, Security Manager warns you when you create an object whose definition is identical to that of an existing object, but it does not prevent you from proceeding. For more information, see Policy Objects Page, page 11-66.

Related Topics

- Chapter 6, "Managing Policy Objects"
- Working with Policy Objects—Basic Procedures, page 6-9

Step 1 Do one of the following:

- Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4. Select the type of object you want to create from the table of contents, right-click in the table and select **New Object**.
- While configuring a rule, click **Select** next to a field that allows or requires a policy object. In the object selector, click the **Create** button below the available objects list.

Tip In a few cases, such as network/host and service objects, clicking these buttons opens a list from which you need to select a specific type for the object.

The dialog box for adding the selected type of object opens. For more information about the individual types of objects, see the following topics:

- Understanding AAA Server and Server Group Objects, page 6-27
- Creating Access Control List Objects, page 6-53
- Add or Edit As Path Object Dialog Boxes, page 56-151
- ASA Group Policies Dialog Box, page 34-1
- Add or Edit BFD Template Dialog Box
- Using Category Objects, page 6-13
- Add or Edit Community List Object Dialog Box, page 56-153
- Configuring Credentials Policy Objects, page 28-9
- Add and Edit File Object Dialog Boxes, page 34-37
- Understanding FlexConfig Policies and Policy Objects, page 7-2 and Creating FlexConfig Policy Objects, page 7-28 (in Chapter 7, "Managing FlexConfigs")
- Creating Identity User Group Objects, page 13-19
- Configuring IKEv1 Proposal Policy Objects, page 26-10

- Configuring IKEv2 Proposal Policy Objects, page 26-14
- Understanding Interface Role Objects, page 6-73
- Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27
- Add and Edit LDAP Attribute Map Dialog Boxes, page 6-46
- Understanding Map Objects, page 6-78
- Understanding Networks/Hosts Objects, page 6-80
- PKI Enrollment Dialog Box, page 26-58
- Add or Edit Policy List Object Dialog Box, page 56-143
- Understanding Pool Objects, page 6-92
- Add or Edit Port Forwarding List Dialog Boxes, page 34-40
- Configuring Port List Objects, page 6-102
- Add or Edit Prefix List Object Dialog Box, page 56-146
- Configuring Risk Rating Policy Objects, page 40-15
- Add or Edit Route Map Object Dialog Boxes, page 56-136
- Creating Security Group Objects, page 14-14
- Creating Cisco Secure Desktop Configuration Objects, page 33-18
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Add or Edit Single Sign On Server Dialog Boxes, page 34-42
- Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8
- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78
- Add or Edit SSL VPN Gateway Dialog Box, page 34-64
- Add and Edit Smart Tunnel Auto Signon List Dialog Boxes, page 34-71
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Add or Edit Text Object Dialog Box, page 7-32
- Configuring Time Range Objects, page 6-71
- Configuring Traffic Flow Objects, page 58-18
- Add or Edit User Group Dialog Box, page 34-73
- Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs, page 31-88

Step 2 Enter a name for the object and optionally a description of the object.

Object names are not case-sensitive and are limited to 128 characters. You can begin object names with a letter, a number, or an underscore. You can use a mix of letters, numbers, special characters, and spaces for the remainder of the object name.

Supported special characters include

- hyphens (-),
- underscores (_),
- periods (.), and,

- plus signs (+).

Beginning with version 4.12, Security Manager allows you to use additional special characters including

- exclamation mark (!),
- at sign (@),
- hash sign (#),
- percent sign (%),
- ampersand sign (&), and,
- parentheses or round brackets ().

Security Manager does not support the following characters:

- caret character (^)
- dollar character ($)

Some objects also support the use of colons (:) in the object name; however, objects with a colon in the name are not supported on IPS devices. If you share objects between different device types that include IPS devices, you should avoid using a colon (:) in the object name.

Note Certain object types, such as AAA server groups, ASA user groups, maps, network/host objects, service objects, and traffic flows, have different naming guidelines. For more details, refer to the online help when you are creating each object type.

Step 3 Configure the settings specific to the type of object. Refer to the online help page for the dialog box.

Step 4 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 5 (Optional) If the object type provides the option, select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 6 Click **OK** to save the object.

Editing Objects

You can edit any user-defined object as required. Changes that you make to the object are reflected in all policies (and other objects) that use the object. However, if an override for the object is already defined for a device, your edits are not reflected in the object used on those devices.

Tips

- You cannot edit predefined objects, but you can copy them to create new objects. See Cloning (Duplicating) Objects, page 6-14.
- Messages appear at the top of the Edit dialog box to indicate the following situations:
  - That you have read-only access to the object. You cannot save changes to these objects.
  - That the policy object was imported using the procedure described in Importing Policies or Devices, page 10-13. Imported objects might be re-imported at some point if the shared policy that uses the object is managed on a different server. Any changes that you make are eliminated if the policy object is imported again. Before editing the object, ensure that you understand the

protocols used in your organization for policy management and importation. You can control whether this message appears using an option on the Tools > Security Manager Administration > Policy Management page (see Policy Management Page, page 11-64).

- You can also edit objects when you define policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Before You Begin

Determine if the object is being used, and which policies, objects, and devices would be affected by the changes. You can generate a usage report for this purpose. See Generating Object Usage Reports, page 6-15.

Related Topics

- Creating Policy Objects, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Step 2 Select the object type from the table of contents.

Step 3 Right-click the object you want to edit and select **Edit Object**.

Step 4 Modify the fields in the Edit dialog box for that object type as required, then click **OK** to save your changes. Click the Help button for information specific to the type of object.

Using Category Objects

Categories provide an intermediate level of detail to objects. By assigning a category to an object, you can look for the name and color of a category to more easily identify rules and objects in rules tables. You can assign a category to a rule or object when you create the rule, or you can edit the rule or object to include category information later. No device configuration commands are generated for category assignments.

The benefits of assigning categories to policy objects are:

- Visibility is improved when you view rules tables using objects that are categorized.
- Objects can be filtered in the rules tables based on category, facilitating rule maintenance.

For example, you might want to create a network/host object and keep track of its use for administrative purposes. When you define this network/host object, you associate it with a category. When you view the access rules table, you can easily identify those rules that use your network/host object. You can also filter the table to display only those items associated with the category.

Security Manager includes a set of predefined categories. Although you cannot change the colors, you can change their names and descriptions. The following procedure explains how to change the name and description.

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **Categories** from the Object Type selector.

Step 3 Click **Edit Object** to open the Category Editor dialog box.

Step 4 Modify the names and descriptions of the predefined category objects as required:

- Label—The color associated with the category.
- Name—The category name. Names can have a maximum of 128 characters, including special characters and spaces.
- Description—Additional information about the object (up to 1024 characters).

Step 5 Click **OK** to save your changes.

Cloning (Duplicating) Objects

An alternative to creating a policy object from scratch is to clone, or duplicate, an existing object. The new object contains all the attributes of the copied object. You can then modify the name and all attributes as required.

Cloning is useful for creating objects that are based on predefined objects that cannot be edited.

Related Topics

- Working with Policy Objects—Basic Procedures, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Step 2 Select the object type from the table of contents.

Step 3 Right-click the object you want to duplicate and select **Clone Object**.

The dialog box for that object type appears. The Name field contains the following default name for the new object: Copy of *name of copied object*. The remaining fields contain the same values as the copied object.

Step 4 Modify the name of the new object and its configuration, as required. Click the Help button for information specific to that type of object.

Step 5 Click **OK** to save your changes.

Viewing Object Details

You can view contents of an object in read-only mode, even when the object is locked by another activity. This is useful when you need to view complete configuration details for complex objects whose definitions cannot be fully displayed in the Policy Object Manager window or when your user privileges allow you only to view object information.

Related Topics

- Working with Policy Objects—Basic Procedures, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Step 2 Select the object type from the table of contents.

Step 3 Right-click the object and select **View Object**.

The dialog box for that object appears in read-only mode.

Generating Object Usage Reports

Before you make any changes to a policy object, you should determine if the object is being used. You can do this by viewing the Referenced column in the Policy Object Manager window. Select the Referenced button above the Policy Object Table to enable the Referenced column.

For objects that are referenced, you can generate usage reports that show which policies, objects, VPNs, and devices are using the selected object and would therefore be affected by changes to that object. Usage reports contain any references to the selected object in your current activity as well as references found in the data committed to the database.

Note The Referenced column reports usage based on both committed data and uncommitted data across all activities/tickets, whereas as the Find Usage feature only reports usage based on committed data and data from the current activity/ticket.

You can use either of these methods to generate usage reports:

- Policy Object Manager—Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4. Select the type of object from the table of contents, right-click the object and select **Find Usage**.
- Firewall rules policies—Left-click an object in a firewall rules table, then right-click and select **Find Usage**.

The usage information is displayed in the Object Usage dialog box. Select the appropriate usage type above the table to view devices, policies, VPNs, or other objects that use the selected object.

For certain policies,

The following table describes the fields in the dialog box.

Table 6-4 Object Usage Dialog Box

Element	Description
Name Type Description	General information about the object for which you are finding usage is displayed at the top of the Object Usage dialog box.
Devices Policies Objects VPN	The type of references you want to view. For example, you can select Objects to view only references to the object from other objects.
Used By	The name of the device, policy, VPN, or object that is referencing the selected object.
Type	The type of item that is referencing the selected object. This can be a device, policy, VPN, or another object.

Table 6-4 Object Usage Dialog Box (continued)

Element	Description
Usage	Indicates how the object is being referenced. For example, if a device is referencing the selected object, this column will indicate that it is a policy assigned to the device that is referencing the object.
Proximity	Indicates the relationship between the selected object and the item that it using it. For example: • A policy that includes a network/host object in its definition has a *direct* relationship with the object and an *indirect* relationship with any other network/host objects contained within the object. • A device on which this policy is assigned references the network/host object *directly* and any other network/host objects contained within the object *indirectly*.
Details Panel	Shows additional details for certain types of references: • **Devices** - For supported policy types, device information is displayed in the Details panel. • **Policies** - For the following supported policy types, the actual rules referencing the object are presented in the Details panel: – AAA Rules – Access Rules – IPv6 Access Rules – Inspection Rules – Translation Rules – Web Filter Rules (PIX/FWSM/ASA) – Zone Based Firewall Rules You can navigate to the rule, export the rule data, or print the rule data from the Details panel. • **Objects** - Details for other objects that are referencing the specified object are presented in the Details panel. You can export the detailed information, print the information, view the object in read-only mode, edit the object, or even find usage for the object from the Details panel in the Object Usage dialog box.

Deleting Objects

You can delete user-defined objects only when they are not being used by policies or other objects. You cannot delete predefined objects. If you delete an object for which device-level overrides are defined, all overrides are also deleted.

Tip You might be prevented from deleting an unused object from the database, if, for example, you replace a local policy that used the object with a shared policy that does not. If object deletion fails, submit or discard all pending changes (in Workflow mode, submit or discard all pending activities), then try again to delete the object. Alternatively, you can leave unused objects in the database, because they will not affect your policies.

Before You Begin

Determine if the object is currently being used and which policies, objects, and devices would be affected by the deletion. You need to remove all references to the object before you can delete it. You can generate a usage report for this purpose. See Generating Object Usage Reports, page 6-15.

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Step 2 Select the object type from the table of contents.

Step 3 Right-click the object you want to delete and select **Delete Object**, or select the object and click the **Delete Object** button. You are asked to confirm the deletion.

Managing Object Overrides

When you create a policy object, you can elect to allow the object to be overridden. This makes it possible to create a generic object to enable you to create general policies. For individual devices, you override the policy object definition to make the policy apply correctly to the device.

From the Policy Object Manager, page 6-4, you can select a policy object that can be overridden and generate a table of device-level overrides that are defined for that global object. Right-click the object and select **Edit Device Overrides** to generate the table (see Policy Object Overrides Window, page 6-20).

You can create device-level overrides in two places:

- In the Device Properties window of a selected device, which allows you to create and manage overrides for the selected device only. For more information, see Creating or Editing Object Overrides for a Single Device, page 6-19.
- In the Policy Object Manager window, which allows you to create and manage overrides for more than one device at a time. For more information, see Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19.

Tip If you override any part of the object definition at the device level, any subsequent changes made to the policy definition at the global level do not affect the device on which the object was overridden.

The following topics explain policy object overrides in more detail:

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19
- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19
- Deleting Device-Level Object Overrides, page 6-21

- Overridable Objects in Security Manager, page 6-22

Understanding Policy Object Overrides for Individual Devices

For many types of policy objects, you can elect to allow an object to be overridden for a particular device. Thus, you can create an object whose definition works for most devices, and then create modifications to the object for the few devices that need slightly different definitions. Or, you can create an object that needs to be overridden for all devices, but which allows you to create a single policy for all devices. Object overrides make it possible for you to create a smaller set of shared policies for use across your devices without giving up the ability to alter policies when needed for individual devices.

For example, you might want to deny ICMP traffic to the different departments in your company, each of which is connected to a different network. You can do this by defining an access rule firewall policy with a rule that includes a network/host object called Departmental Network. By allowing device override for this object, you can then create overrides on each relevant device that specify the actual network to which that device is connected.

Device-level object overrides are especially important when the global object is included in the definition of a VPN policy, which applies to every device in the VPN topology. For example, you select a PKI enrollment object when defining a PKI policy on a site-to-site VPN. If the hub of your VPN uses a different CA server than the spokes, you must use device-level overrides to specify the CA server used by the hub. Although the PKI policy references a single PKI enrollment object, the actual CA server represented by this object will differ for the hub, based on the device-level override you define.

You can quickly tell if an object can be overridden by looking for the Overrides column in the objects table in the Policy Object Manager, page 6-4. A green checkmark indicates that you can create overrides for the object; the presence of the column indicates the object type allows overrides.

Related Topics

- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19
- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19
- Deleting Device-Level Object Overrides, page 6-21

Allowing a Policy Object to Be Overridden

To create overrides for an object, the object must allow overrides. Not all object types allow overrides.

For those that do allow overrides, you define the object as allowing overrides by selecting **Allow Value Override per Device** when defining the object. After selecting this option, you must click **OK** to save the object before you can define any overrides. For more information on creating objects, see Creating Policy Objects, page 6-9.

You can also configure Security Manager to create device-level overrides for existing objects when you discover policies on devices that you add to the inventory. During discovery, if Security Manager determines that an existing object applies to a discovered policy, but that it is not a perfect fit, the object is used but a device-level override is created to account for the difference. For example, if you run policy discovery on a device that has an ACL with the same name as an ACL policy object in Security Manager, the name of the discovered policy object is reused, but a device-level override is created for the object. If you do not allow device-level overrides during discovery, a new policy object is created with a number appended to the name; this is the default.

To configure Security Manager to allow device overrides during discovery, select **Tools > Security Manager Administration > Discovery** and select **Allow Device Override for Discovered Policy Objects**.

Note To ensure that a specific policy object will be reused for device-level override during discovery, make sure the **Allow Value Override per Device** check box has been selected for the policy object in Policy Object Manager before policy discovery.

Related Topics

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19
- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19
- Deleting Device-Level Object Overrides, page 6-21

Creating or Editing Object Overrides for a Single Device

You can create or edit device-level object overrides from the Device Properties window.

An override specifies a definition for a global object that affects only the selected device. For example, you can override the definition of a AAA server group object so that the object represents a different group of AAA servers for one device than the group it represents for other devices.

Related Topics

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19
- Deleting Device-Level Object Overrides, page 6-21

Step 1 (Device view) Right-click a device in the Device selector and select **Device Properties**.

Step 2 Select the object type you want to override from the **Policy Object Overrides** folder.

The table displays all objects of the selected type that can be overridden at the device level. If an object has an override already defined for the device, the Value Overridden? column contains a check mark.

Step 3 Select the object whose override you want to change and do one of the following:

- Click the **Create Override** button, or right-click and select **Create Override**.
- Click the **Edit Override** button, or right-click and select **Edit Override**.

The dialog box for defining that type of object is displayed with the current properties (either the global properties or the local override).

Step 4 Modify the definition of the object and click **OK** to save the device-level override. In the Device Properties window, a green check mark appears in the Value Overridden? column.

Creating or Editing Object Overrides for Multiple Devices At A Time

You can create or edit device-level object overrides from the Policy Object Manager window.

This method enables you to create overrides on multiple devices at the same time, which is especially useful when creating overrides for several devices that participate in the same VPN topology. For example, if the spokes located in one part of the VPN use a different CA server than the spokes located in a different part of the VPN, you can override the PKI enrollment object that defines the server for these devices. This is a more convenient method than selecting each device individually from Device view and defining the override from the Device Properties window.

Related Topics

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19
- Deleting Device-Level Object Overrides, page 6-21

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Step 2 Select the object type you want to override from the table of contents, and then select the object to override.

Tip Not all types of object allow overrides, and not all objects are defined as overridable. Look for a green check mark in the Overridable column. If the object type allows overrides, but this object does not have a check mark, edit the object to enable object override (see Allowing a Policy Object to Be Overridden, page 6-18).

Step 3 Double-click the checkmark, or right-click the object and select **Edit Device Overrides**, to open the Policy Object Overrides Window, page 6-20. The window contains a table listing each device for which an override is defined for the object.

Tip You can also edit the overridable object and click **Edit** next to the Overrides field.

Step 4 Do one of the following:

- To add an override, click the **Create Override** button, select the devices to which you want to apply the override, and define the override.

  The dialog boxes for creating and editing the override are the same ones used to create the object; click the Help button for information specific to the type of object.

  The override you create applies to all policies on the device that use the object; you cannot override the object for one policy but not for another policy.
- To edit an override, select it and click the **Edit Override** button.

Policy Object Overrides Window

Use the Policy Object Overrides window to view a list of all device-level overrides that are defined for the selected object. The content displayed in the table differs depending on the type of object, but it always includes the device name, object description, and category. Sometimes the content of the object is shown, including the overrides.

- To add an override, click the **Create Override** button. In the Create Overrides for Device window, select the devices from the available list and click >> to move them to the selected list. When you click **OK**, you are presented with the dialog box for defining your override, which applies to all newly selected devices. (You are not changing the override of the greyed out devices)

  Note The available devices list shows the devices that have not already had overrides defined for the object. Devices with overrides are shown greyed out in the selected devices list.

  The dialog boxes for creating and editing the override are the same ones used to create the object; click the Help button for information specific to the type of object.

  The override you create applies to all policies on the device that use the object; you cannot override the object for one policy but not for another policy.

- To edit an override, select it and click the **Edit Override** button.
- To delete an override, select it and click the **Delete Override** button.

  Deleting an override does not delete the object or remove the object from its device assignment. When you delete the override, the policies on the device that use the object start using the global definition for the object. This changes the meaning of the policies.

Tip You can also create and edit device-level overrides from the Device Properties window of a selected device. Using the Device Properties windows makes it easy for you to manage the overrides for all objects used by a single device. For more information, see Creating or Editing Object Overrides for a Single Device, page 6-19.

Navigation Path

Open the Policy Object Manager, page 6-4. Select an object type that can be overridden (its object page contains a column called Overrides), then do one of the following:

- Double-click the green checkmark in the Overrides column.
- Right-click the object and select **Edit Device Overrides**.
- Edit the overridable object and click **Edit** next to the Overrides field.

Related Topics

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19
- Deleting Device-Level Object Overrides, page 6-21
- Filtering Tables, page 1-48
- Filtering Items in Selectors, page 1-45

Deleting Device-Level Object Overrides

Deleting a device-level override restores the global definition of the object to the selected device. You can delete overrides from the Device Properties window or from the Policy Object Manager window:

- Deleting overrides from Device view—Right-click the device and select **Device Properties**, then select the object type from the **Policy Object Overrides** folder. Select the override you want to delete and click **Delete Override**.
- Deleting overrides from the Policy Object Manager—Select the object type from the table of contents, then right-click the object and select **Edit Device Overrides**. Select the override you want to delete and click **Delete Override**.

Related Topics

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Policy Object Override Pages, page 3-52
- Policy Object Overrides Window, page 6-20

Overridable Objects in Security Manager

You can override the following objects in Security Manager:

- **VPN Objects**
  - AAA Server group
  - PKI Enrollment (CA Servcers)
  - WINS Server List
  - SSL VPN Customization
  - SAML Identity Provider
  - Web ACL
  - Port Forwarding List
  - Bookmarks
  - Smart Tunnel List
  - Smart Tunnel Network List
  - Smart Tunnel Auto Sign on List
  - Single Sign on Server
  - Reference Identity
- **Firewall Objects**
  - Identity User Group
  - Networks/Hosts
  - Port Lists
  - Security Group
  - Services
  - Access Control Lists (Extended, Standard, Web, Unified)
  - As Path
  - BFD Template
  - Community List

- Credentials
- Identity Policy (IOS)
- Identity User Group
- Interface Roles
- LDAP Attribute Maps
- LDAP Attribute Maps (IOS)
- Policy List
- Prefix List
- Prefix Lists IPV6
- Risk Rating
- Route Map
- Security Group
- Text Objects
- TLS Proxy
- Pool Objects (DHCP V6,IPV4 Pool,IPV6 Pool, MAC Address Pool, NET Pool)
- MAPs (AVP, Regular Expression Groups, Regular Expressions, TCP Maps)
- Class Maps—Inspect (AOL,DCE/RPC,DIAMETER,DNS,eDonkey,FastTrack,FTP,GunTella,H.323(ASA/PIX/FWSM),H.323(IOS), HTTP(ASA/PIX/FWSM),HTTP(IOS),ICQ,IM,IMAP,Kazaa2,MSN Messenger,POP3,Scansafe,SIP(ASA/PIX/FWSM), SIP(IOS),SMTP,SUN RPC, Windows Messenger, Yahoo Messenger)
- Class Maps—Web Filter (Local,N2H2,Trend,Websense)
- Parameter Maps—Inspect (Inspect Parameters, Protocol Info Parameters)
- Parameter Maps—Web Filter(Loal,N2H2,Trend,URL Filter, URLF Glob parameters, Websense)
- Policy Maps—Inspect (DCE/RPC,DIAMETER,DNS,ESMTP,FTP,GTP,H.323(ASA/PIX/FWSM),H.323(IOS),HTTP ASA7.1.x/PIX7.1.x/FWSM3.x/IOS), HTTP(ASA7.2+/PIX7.2+),HTTP(Zone Based IOS),IM(ASA7.2+/PIX7.2+),IM(IOS),IM(Zone Based IOS),IMAP, IP Options, IPSec Pass Trough,IPV6,LISP,M3UA,NetBIOS,P2P,POP3,Scansafe,Sctp,SIP(ASA/PIX/FWSM), SIP(IOS), Skinny, SMTP, SNMP,SUN RPC)
- Policy Maps—Web Filter (Web Filter)

Importing and Exporting Policy Objects

Security Manager includes a Perl script that you can use to export network/host, service, and port list policy objects so that you can import them into another Security Manager server. The information includes device-level overrides for policy objects that have them.

Note The command works with network/host objects that contain IPv4 addresses only. You cannot use the command to import network/host-IPv6 objects.

You can also manually create a CSV file that you can import. For example, you might obtain a list of IP addresses that identify networks or hosts that should be denied entry to your network. You can create a CSV file that will bulk-load the list as one or more network/host objects if that is easier than manually creating the object in the Policy Object Manager.

Tip Besides using this command, you can use other facilities to export and import policy objects that are assigned to shared policies or configured in local device policies. For more information, see the following topics: Exporting the Device Inventory from the Security Manager Client, page 10-6, Exporting Shared Policies, page 10-12, and Importing Policies or Devices, page 10-13

The Perl command is located in $NMSROOT\bin, which is typically C:\Program Files\CSCSpx\bin. The syntax of the command is:

perl [*path*]**PolicyObjectImportExport.pl -u** *username* **-p** *password* **-o {import | export}** [**-a** *activity*] **-t** *object_type* **-f** *filename* [**-c {true | false}**] [**-d {true | false}**] [**-e {true | false}**] [**-g {true | false}**] [**-h**]

Syntax

perl [*path*] **PolicyObjectImportExport .pl**	The Perl script command. Include the path to the PolicyObjectImportExport.pl file if the path is not defined in the system path variable. **Tip** If you forget to include the "perl" command, the system accepts the input but does nothing and provides no feedback on your error. Use Ctrl+Z to return to the command prompt.
-u *username*	A Security Manager username. The data exported is limited by the permissions assigned to this user. The user must have Modify Objects permission for the import or export of policy objects, and additionally the Modify Devices permission for the import or export of device-level overrides. If you are importing objects in non-Workflow mode, you must also have Submit and Approve privileges.
-p *password*	The user's password.
-o {import | export}	The type of operation you are performing, either to import policy objects from an existing file, or to export policy objects to a CSV file. Only committed objects are exported.
-a *activity*	(Optional.) The name of a Workflow activity. If you do not specify a name, a new activity is created with the name username_time.
-t *object_type*	Object type, one of the following: • network—For network/host objects. • service—For service objects. • portlist—For port-list objects.
-f *filename*	The name of the CSV file. When exporting, if the file exists, it is overwritten.

-c {true | false}	(Optional.) When importing objects, whether to enable policy object conflict detection. • false—An object is imported even if an existing object has the same content. • true—If an existing object has the same content as an imported object, the imported object is skipped. You must also select **Enforce** for the **When Redundant Objects Detected** option on the Policy Objects Page, page 11-66.
-d {true | false}	(Optional.) How to handle device-level policy object overrides during either an import or export operation: • true—Include all globally-defined objects and all device-level overrides of the objects. • false—Include only the global definitions of the policy objects. Do not include any device-level policy object override information. This is the default.
-e {true | false}	(Optional.) Whether to "flatten" port-list objects in service objects and service-group objects: • true—The names of any port-list objects found in service objects and service-group objects are replaced with the actual ports from the lists. That is, the two objects, port-list and service, or port-list and service-group, are "flattened" into a single service or service-group. Port-list objects are used in Security Manager to group sets of port definitions, and are used when defining service and service-group objects. However, port-list objects are not supported in PRSM. • false—Port-list objects in service and service-group objects are not flattened. This is the default.
-g {true | false}	(Optional.) Whether to include object and object-group types in the CSV file: • true—The final column in the file will be **Type** and it will indicate "Service" or "Network." • false—The Type column is not included. This is the default.
-h	(Optional.) Display the command line help. If you include this option, all other options are ignored.

Importing Policy Objects

When you are importing objects, if an object refers to another object, that object must already be defined in Security Manager, or it must be defined in the same CSV file that you are importing. If the object is in the same CSV file, it must come before the object that refers to it. (Security Manager automatically sorts objects as required when exporting them.)

If Security Manager already has a policy object of the same name as one you are importing, the object is skipped and not imported. The name conflict can even occur if another user has created an object but not yet committed it for public viewing, so you might not be able to see the conflicting object. Security Manager creates only new objects, it does not update existing objects. Use the -c option to specify whether new objects can be created that have the same content as existing objects.

When you run the command, if there are any errors in the file, only the affected objects are not imported. Error messages indicate these problems as they occur, and Security Manager continues evaluating all records in the file. All correctly defined policy objects are imported, and the objects with errors are skipped. The total count and the names of the policy objects that are not imported are shown in the output screen.

After the import command completes, additional actions depend on the Workflow mode you are using:

- Workflow mode—You must log into Security Manager using the same username and password and submit the activity you specified during the import. The activity must be submitted and approved for the changes to take effect.
- Non-Workflow mode—The imported objects are automatically submitted and approved without action on your part. However, you will receive an error if the username you supplied does not have Submit and Approve privileges, and the import operation will fail.

CSV File Format

All objects in a single file are of the same policy object type. The file is in standard comma-separated values (CSV) format. The first line has column headings. Each row represents a single policy object. The columns, left to right, are:

- Name—(Mandatory.) The name of the object.
- Node—The display name of the device on which an override of the policy object is defined. If the policy object is defined on the global level, the field is empty. When importing objects, if the display name does not match a device already in the Security Manager inventory, the object is skipped and not imported.
- Description—The description of the object, if any.
- Category—The category identifier of the object, if any. The category ID is from 10 to 19.
- Allow Override—Whether the object can be overridden. True if the policy object can be overridden on device level, False (or an empty field) if not.
- Group—The names of other policy objects with the same type referenced by this policy object. If there is more than one object, they are separated by commas. For example, network building block Net1 references network building block Net2 and Net3. The Group field of Net1 would have “Net2,Net3” as its value.
- Data—The content of the object.
- Subtype—The object subtype, if any, for network/host and service objects. For an explanation of network/host and service object types, see Understanding Networks/Hosts Objects, page 6-80 and Understanding and Specifying Services and Service and Port List Objects, page 6-100. Possible values are:
  - Blank, or space—The object is a group object, either network/host or service.
  - NH—(Network/host objects only.) Single host network/host object.
  - NF—(Network/host objects only.) Single fully-qualified domain name (FQDN) network/host object.
  - NN—(Network/host objects only.) Single network address network/host object.
  - NR—(Network/host objects only.) Single Address range network/host object.
  - SO—(Service objects only.) Single-service service object.
- Type—The type of object represented by this entry: “Network” or “Service.”

If there is no value for a particular field, that field is blank in the output. If there are multiple values for a field, the field is enclosed in double quotation marks.

Understanding AAA Server and Server Group Objects

You use AAA server objects to identify the AAA servers used in your network. AAA enables devices to determine who the user is (authentication), what the user is permitted to do (authorization), and what the user actually did (accounting), as described below:

- Authentication—Authentication is the way a user is identified before being allowed access to the network and network services. It controls access by requiring valid user credentials, which are typically a username and password. All authentication methods, except for local, line password, and enable authentication, must be defined through AAA. You can use authentication alone or with authorization and accounting.
- Authorization—After authentication is complete, authorization controls the services and commands available to each authenticated user. Authorization works by assembling a set of attributes that describe what the user is authorized to perform. These attributes are compared to the information contained in a database for a given user and the result is returned to AAA to determine the user's actual capabilities and restrictions. The database can be located locally on the access server or router or it can be hosted remotely on a RADIUS or TACACS+ security server. Were you not to use authorization, authentication alone would provide the same access to services to all authenticated users. You must use authorization together with authentication.
- Accounting—Accounting is used to track the services users are accessing, as well as the amount of network resources they are consuming. When AAA accounting is activated, the network access server reports user activity to the RADIUS or TACACS+ security server (depending on which security method you have implemented) in the form of accounting records. Accounting information includes when sessions start and stop, usernames, the number of bytes that pass through the device for each session, the service used, and the duration of each session. This data can then be analyzed for network management, client billing, or auditing. You can use accounting alone or together with authentication and authorization.

AAA provides an extra level of protection and control for user access over using access rules (ACLs) alone. For example, you can create an access rule allowing all outside users to attempt to use Telnet on a server on the DMZ network. If you want only some users to actually reach the server (and you might not always know the IP addresses of these users, making it impossible to configure simple access rules), you can enable AAA to allow only authenticated or authorized users to make it through the network device (for example, the ASA or router). Thus, users must authenticate before reaching the Telnet server, where Telnet can also require a separate login.

AAA server objects are collected into AAA server group objects. Policies requiring AAA (such as Easy VPN, Remote Access VPNs, and router platform policies such as Secured Device Provisioning and 802.1x) usually refer to AAA server group objects. These objects contain multiple AAA servers that use the same protocol, such as RADIUS or TACACS+. In essence, AAA server groups represent collections of authentication servers focused on enforcing specific aspects of your overall network security policy. For example, you can group those servers dedicated to authenticating internal traffic, external traffic, or remote dial-in users, as well as servers that authorize the administration of your firewall devices.

The following topics describe how to work with AAA server objects:

- Supported AAA Server Types, page 6-28
- Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28
- Predefined AAA Authentication Server Groups, page 6-30
- Default AAA Server Groups and IOS Devices, page 6-31
- Creating AAA Server Objects, page 6-32
- Add or Edit AAA Server Dialog Box, page 6-33

- Add and Edit LDAP Attribute Map Dialog Boxes, page 6-46
- Creating AAA Server Group Objects, page 6-48

Supported AAA Server Types

You can use AAA servers that use the RADIUS protocol with all devices, and the TACACS+ and LDAP protocols with all devices except IPS. For ASA, PIX, and FWSM devices, you can also use the protocols described in Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28

- **RADIUS**—Remote Authentication Dial-In User Service (RADIUS) is a distributed client/server system that secures networks against unauthorized access. In the Cisco implementation, RADIUS clients run on Cisco devices and send authentication requests to a central RADIUS server that contains all user authentication and network service access information.

  You can use RADIUS with other AAA security protocols, such as TACACS+, Kerberos, and local username lookup, depending on what is supported by a particular device type. RADIUS is supported on all Cisco platforms, but some RADIUS-supported features run only on specified platforms.

  Beginning with Cisco Security Manager 4.17, IPv6 is enabled in RADIUS protocol. This support is applicable only for ASA 9.9.2 devices and above. Users can now configure IPv6 Host Address for Radius authentication in the Add AAA Server dialog box (see, Add or Edit AAA Server Dialog Box). Activity validation is also introduced for unsupported device version.
- **TACACS+**—Terminal Access Controller Access Control System (TACACS+) is a security application that provides centralized validation of users attempting to gain access to a router or network access server. The goal of TACACS+ is to provide a methodology for managing multiple network access points from a single management service.

  TACACS+ provides for separate and modular authentication, authorization, and accounting facilities. TACACS+ allows for a single access control server (the TACACS+ daemon) to provide each service independently.
- **LDAP**—Lightweight Directory Access Protocol (LDAP). The use of LDAP servers is specific to certain policies. For example, identity firewall configurations on ASA, VPN configurations on ASA, and ScanSafe configurations on IOS devices. For more information on using LDAP on ASA, see Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28.

Related Topics

- Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28
- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27

Additional AAA Support on ASA, PIX, and FWSM Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

In addition to supporting RADIUS and TACACS+, ASA, PIX 7.0+, and FWSM 3.1+ devices can support AAA servers running the following protocols. For more information, see the explanation of AAA usage in the configuration guides for the device type and operating system version that interests you.

- **Kerberos**—These devices can use Kerberos servers for authentication. 3DES, DES, and RC4 encryption types are supported.
- **NT**—These devices can use Windows Domain servers for NTLMv1 authentication.
- **SDI Servers**—SecureID servers from RSA Security, Inc. are known as SDI servers. When a user attempts to establish VPN access and the applicable tunnel-group policy specifies an SDI authentication server group, the ASA device sends the username and one-time password to the SDI server. The device then grants or denies user access based on the response from the server. Version 5.0 of SDI introduced the concept of SDI primary and secondary servers that share a single-node secret file (SECURID). As a result, when you configure an SDI server as a AAA server object, you must specify whether the server is version 5.0 or an earlier version.
- **LDAP**—These devices can use Lightweight Directory Access Protocol (LDAP) servers for VPN authorization and user group identification for identity-aware firewall policies. These devices support LDAP version 3 and are compatible with any v3 or v2 directory server. However, password management is supported only on the Sun Microsystems JAVA System Directory Server and the Microsoft Active Directory.

  With any other type of LDAP server (such as Novell or OpenLDAP), all LDAP functions are supported except for password management. Therefore, if someone tries to log in to one of these devices using one of these other servers for authentication and their password has expired, the device drops the connection and a manual password reset is required.

  You can configure Simple Authentication and Security Layer (SASL) mechanisms to authenticate an LDAP client (in this case, the ASA, PIX, or FWSM device) to an LDAP server. These devices and LDAP servers can support multiple mechanisms. If both mechanisms (MD5 and Kerberos) are available, the ASA, PIX, or FWSM device uses the stronger mechanism, Kerberos, for authentication.

  When user authentication for VPN access has succeeded and the applicable tunnel-group policy specifies an LDAP authorization server group, the ASA, PIX, or FWSM device queries the LDAP server and applies the authorizations it receives to the VPN session.
- **HTTP-Form**—These devices can use the HTTP Form protocol for single sign-on (SSO) authentication of WebVPN users only. Single sign-on support lets WebVPN users enter a username and password only once to access multiple protected services and Web servers. The WebVPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the WebVPN server sends an SSO authentication request, including username and password, to the authenticating server using HTTPS. If the server approves the authentication request, it returns an SSO authentication cookie to the WebVPN server. The security appliance keeps this cookie on behalf of the user and uses it to authenticate the user to secure websites within the domain protected by the SSO server.

The following table describes the AAA services that are supported by each protocol:

Table 6-5 AAA Services Supported by ASA, PIX, and FWSM Devices

	Database Type							
AAA Service	**Local**	**RADIUS**	**TACACS +**	**SDI**	**NT**	**Kerberos**	**LDAP**	**HTTP Form**
Authentication of...								
VPN users	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes[1]
Firewall sessions	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Administrators	Yes	Yes	Yes	Yes[2]	Yes	Yes	Yes	No

Table 6-5 AAA Services Supported by ASA, PIX, and FWSM Devices (continued)

AAA Service	Database Type							
	Local	RADIUS	TACACS+	SDI	NT	Kerberos	LDAP	HTTP Form
Authorization of...								
VPN users	Yes	Yes	No	No	No	No	Yes	No
Firewall sessions	No	Yes[3]	Yes	No	No	No	No	No
Administrators	Yes[4]	No	Yes	No	No	No	No	No
Accounting of...								
VPN connections	No	Yes	Yes	No	No	No	No	No
Firewall sessions	No	Yes	Yes	No	No	No	No	No
Administrators	No	Yes[5]	Yes	No	No	No	No	No

1. HTTP Form protocol supports single sign-on (SSO) authentication for WebVPN users only.
2. SDI is not supported for HTTP administrative access.
3. For firewall sessions, RADIUS authorization is supported with user-specific ACLs only, which are received or specified in a RADIUS authentication response.
4. Local command authorization is supported by privilege level only.
5. Command accounting is available for TACACS+ only.

Related Topics

- Supported AAA Server Types, page 6-28
- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27

Predefined AAA Authentication Server Groups

There are several predefined AAA server groups that define an authentication method without specifying particular AAA servers. In policies such as IPSec proposals, you can use these predefined server groups to define the types of AAA authentication to perform and the order in which to perform them.

Table 6-6 on page 6-30 describes the predefined AAA authentication server groups.

Table 6-6 Predefined AAA Authentication Server Groups

Name	Description
Enable	Uses the enable password defined on the device for authentication.
KRB5 KRB5-Telnet	Uses Kerberos 5 for authentication. Use KRB5-Telnet when using Telnet to connect. For Cisco IOS routers, you can use Kerberos 5 client configuration only on selected platforms running IOS Software versions that support this protocol. Server configuration is not supported. The device must include an Advanced series feature set (k9 crypto image).
If-Authenticated	Uses the if-authenticated method, which allows the user to access the requested function if the user is authenticated.

Table 6-6 Predefined AAA Authentication Server Groups (continued)

Name	Description
Line	Uses the line password defined on the device for authentication.
Local Local-case	Uses the local username database (defined on the device) for authentication. Use Local-case if you want the login to be case-sensitive.
None	Uses no authentication.
RADIUS TACACS+	Use RADIUS or TACACS+ authentication. (Does not apply to Cisco IOS routers.) These AAA server groups do not contain any AAA servers. To use one of them when defining a policy, you must create a device-level override and define the AAA servers to associate with the group. For more information, see Creating or Editing Object Overrides for a Single Device, page 6-19.

Related Topics

- Creating AAA Server Group Objects, page 6-48
- Default AAA Server Groups and IOS Devices, page 6-31
- Understanding AAA Server and Server Group Objects, page 6-27

Default AAA Server Groups and IOS Devices

IOS software enables you to define AAA servers either as members of AAA server groups or as individual servers. Security Manager, however, requires all AAA servers to belong to a AAA server group.

Therefore, when you discover an IOS device whose device configuration contains individual AAA servers that do not belong to a AAA server group, Security Manager creates the following server groups to contain these servers:

- For RADIUS: CSM-rad-grp
- For TACACS+: CSM-tac-grp

Both of these special AAA server groups are marked in the Policy Object Manager as the default groups for their protocol. This is indicated by the **Make this Group the Default AAA Server Group** check box.

These groups are created solely for the purpose of management by Security Manager. During deployment, the AAA servers in these special groups are deployed back to the IOS device as individual servers, *not* as part of the group.

You can also create your own default group. The default group can be used in most cases, except when you need to configure multiple AAA server groups that use the same protocol. For example, you might want to define multiple RADIUS groups so that one group can be used for authentication and another group for authorization. Service providers may want to define multiple groups with the same protocol in order to provide customer separation when using VRF.

Note If you use one of these default AAA server groups in a policy defined for a PIX/ASA/FWSM device, the AAA servers are deployed as a group to that device, not as individual servers. This is because all AAA servers on PIX/ASA/FWSM devices must belong to a AAA server group.

Caution We recommend that you use caution when using these default AAA server groups in a policy definition. There are certain commands (for example, **ip radius** and **ip tacacs**, which are configured using the Interface field in the AAA Server dialog box) that can be defined once for each AAA server group and once for all individual AAA servers. Because the AAA servers in the default group are deployed to IOS devices as individual servers, you might inadvertently change the **ip radius** or **ip tacacs** settings for all the individual AAA severs configured on the device, including servers that are not being managed by Security Manager (and whose configurations would otherwise be left undisturbed).

Related Topics

- Predefined AAA Authentication Server Groups, page 6-30
- Creating AAA Server Group Objects, page 6-48
- Understanding AAA Server and Server Group Objects, page 6-27

Creating AAA Server Objects

You can create AAA server objects to populate the AAA server group objects that are referenced by policies such as AAA rules, Easy VPN, and 802.1x. In some cases, AAA server objects are used directly by a policy, such as in AAA policies on IPS devices.

When creating a AAA server object, you must specify the IP address or DNS name of the external AAA server and the protocol used by the server. The other settings required depend on the protocol.

Note On PIX/ASA/FWSM devices, AAA objects in a device configuration that are not referenced by any policies are removed from the device during the next deployment. However, the predefined AAA objects named RADIUS and TACACS+ are never removed from PIX 6.3 devices, even if they are not referenced by any policies.

Related Topics

- Creating Policy Objects, page 6-9
- Supported AAA Server Types, page 6-28
- Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28
- Understanding AAA Server and Server Group Objects, page 6-27

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **AAA Servers** from the Object Type selector.

Step 3 Right-click in the work area, then select **New Object** to open the Add or Edit AAA Server Dialog Box, page 6-33.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Identify the AAA server:

- In the Host field, enter the IP address or for ASA or PIX 7.2+ devices, the host name of the AAA server. You can also enter the name of a network/host object that contains the host IP address, or click **Select** to select the object.

- Optionally, in the Interfaces field, enter the name of an interface or an interface role (which must resolve to a single interface name on the device) whose IP address should be used for all outgoing RADIUS or TACACS+ packets. Do not specify an interface for objects used on an IPS device.
- Optionally, enter the amount of time to wait until a AAA server is considered unresponsive.

Step 6 Select the protocol used by the AAA server and configure protocol-specific properties. You can use RADIUS with all device types, and TACACS+ with all device types except for IPS devices. You can use the Kerberos, LDAP, NT, SDI, and HTTP-FORM protocols only with ASA, PIX 7.x+, and FWSM 3.1+ devices.

For details about the properties, see the following topics:

- RADIUS—See AAA Server Dialog Box—RADIUS Settings, page 6-35.
- TACACS+—See AAA Server Dialog Box—TACACS+ Settings, page 6-38.
- Kerberos—See AAA Server Dialog Box—Kerberos Settings, page 6-39.
- LDAP—See AAA Server Dialog Box—LDAP Settings, page 6-40.
- NT—See AAA Server Dialog Box—NT Settings, page 6-43.
- SDI—See AAA Server Dialog Box—SDI Settings, page 6-43.
- HTTP-FORM—See AAA Server Dialog Box—HTTP-FORM Settings, page 6-44.

Step 7 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 8 Click **OK** to save the object.

Add or Edit AAA Server Dialog Box

Use Add or Edit AAA Server dialog box to create, copy, and edit a AAA server object. These objects are collected into AAA server group objects and identify the AAA servers that you want to use when defining various AAA policies. In some cases these objects are used directly in a AAA policy.

For a description of the protocols you can use, see Supported AAA Server Types, page 6-28 and Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28.

Note You cannot edit the protocol if the object is already included in a AAA server group.

Navigation Path

Select **Manage > Policy Objects**, then select **AAA Servers** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- Policy Object Manager, page 6-4

Field Reference

Table 6-7 AAA Server Dialog Box—General Settings

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Host	The address of the AAA server to which authentication requests will be sent. Specify one of the following: • IP Address—The IPv4 or IPv6 address of the AAA server. You can also enter the name of a network/host object that contains the host IP address, or click **Select** to select the object. **Note** AAA- IPV6 hosts are only supported for the LDAP and TACACS+ protocols. Beginning with Cisco Security Manager 4.17, IPv6 hosts for the Radius protocol are supported on ASA 9.9(2) devices onwards. • DNS Name (for PIX/ASA 7.2+ devices only)—The DNS hostname of the AAA server, up to 128 characters. The hostname can contain alphanumeric characters and hyphens, but each element of the hostname must begin and end with an alphanumeric character.
Interface	The interface whose IP address should be used for all outgoing RADIUS or TACACS packets (known as the source interface). Enter the name of an interface or interface role, or click **Select** to select it from a list or to create a new interface role. **Tips** • If you enter the name of an interface, make sure the policy that uses this AAA object is assigned to a device containing an interface with this name. • If you enter the name of an interface role, make sure the role represents a single interface, not multiple interfaces. • Only one source interface can be defined for the AAA servers in a AAA server group. An error is displayed when you submit your changes if different AAA servers in the group use different source interfaces. See Creating AAA Server Group Objects, page 6-48. • You cannot specify an interface name for a AAA server used on an IPS device.

Table 6-7 AAA Server Dialog Box—General Settings (continued)

Element	Description
Timeout	The amount of time to wait for a response to a request until the AAA server is considered unresponsive. If there are other servers in the group, the next server is tried. • Cisco IOS routers—The range is 1-1000 seconds. The default is 5 seconds. • ASA/PIX 7.x+, FWSM 3.1+ devices—The range is 1-300 seconds. The default is 10 seconds. • PIX 6.3 firewalls—The range is 1-512 seconds. The default is 5 seconds. • IPS devices—The range is 1-512 seconds. The default is 3 seconds.
Protocol	The protocol used by the AAA server. The fields below the protocol list change depending on your selection. For specific information about the fields, see the topics indicated. • The following protocols are the most common: – RADIUS—All device types. See AAA Server Dialog Box—RADIUS Settings, page 6-35. – TACACS+—All device types except IPS. See AAA Server Dialog Box—TACACS+ Settings, page 6-38. • The following protocols are supported for ASA/PIX 7.x+ and FWSM 3.1+ devices; LDAP is supported on IOS devices that support ScanSafe policies: – Kerberos—See AAA Server Dialog Box—Kerberos Settings, page 6-39. – LDAP—See AAA Server Dialog Box—LDAP Settings, page 6-40. – NT—See AAA Server Dialog Box—NT Settings, page 6-43. – SDI—See AAA Server Dialog Box—SDI Settings, page 6-43. – HTTP-FORM—See AAA Server Dialog Box—HTTP-FORM Settings, page 6-44.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

AAA Server Dialog Box—RADIUS Settings

Use the RADIUS settings in the AAA Server dialog box to configure a RADIUS AAA server object.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **RADIUS** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32

- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-8 AAA Server Dialog Box—RADIUS Settings

Element	Description
Key Confirm	The shared secret that is used to encrypt data between the network device (client) and AAA server. The key is a case-sensitive, alphanumeric string of up to 127 characters. Special characters are permitted. The key you define in this field must match the key on the RADIUS server. Enter the key again in the Confirm field. Note the following: • A key is required for AAA server objects used in an IPS AAA policy. Otherwise, the key is optional. • Spaces are not permitted on PIX, ASA, or FWSM devices. Otherwise, they are permitted. • If you do not define a key, all traffic between the AAA server and its AAA clients is sent unencrypted.
Authentication/Authorization Port	The port on which AAA authentication and authorization are performed. The default is 1645. **Tip** The default port for IPS devices is 1812, so you need to change this value if you are configuring the object for IPS and you want to use the default port.
Accounting Port	The port on which AAA accounting is performed. The default is 1646.

Table 6-8 AAA Server Dialog Box—RADIUS Settings (continued)

Element	Description
RADIUS Password Confirm (ASA, PIX 7.x+, and FWSM 3.x+ devices only.)	A case-sensitive, alphanumeric keyword of up to 127 characters that is common among users who access this RADIUS authorization server through this device. Enter the password again in the Confirm field. The RADIUS authorization server requires a password and username for each connecting user. The RADIUS server administrator must configure the RADIUS server to associate this password with each user authorizing to the server through this device. Be sure to provide this information to your RADIUS server administrator. If you do not specify a common user password, each user password is the username. Never use a RADIUS authorization server for authentication. Common passwords or usernames as passwords are less secure than assigning unique user passwords. **Tips** • The password applies to authorization servers only, not to authentication servers. For an authentication RADIUS servers, do not configure a common password. • Although the password is required by the RADIUS protocol and the RADIUS server for authorization, users do not need to know it. The device provides the password automatically.
Retry Interval (ASA, PIX 7.x+, and FWSM 3.x+ devices only.)	The interval between attempts to contact the AAA server. Values are: • ASA/FWSM devices—1 to 10 seconds. • PIX devices—1 to 5 seconds.

Table 6-8 AAA Server Dialog Box—RADIUS Settings (continued)

Element	Description
ACL Netmask Convert (ASA, PIX 7.x+, and FWSM 3.x+ devices only.)	The method for handling the netmask expressions that are contained in downloadable ACLs received from the RADIUS server. The ASA/PIX/FWSM expects downloadable ACLs to contain standard netmask expressions whereas devices using Cisco IOS Software expect downloadable ACLs to contain wildcard netmask expressions, which are the reverse of a standard netmask expression. A wildcard mask has ones in bit positions to ignore, zeros in bit positions to match. Translation of wildcard netmask expressions means that downloadable ACLs written for Cisco IOS routers can be used by ASA/PIX/FWSM devices without altering the configuration of the ACLs on the RADIUS server. Select one of the following options: • Standard—The security appliance assumes that all downloadable ACLs received from the RADIUS server contain only standard netmask expressions. No translation from wildcard netmask expressions is performed. This is the default. • Auto-Detect—The security appliance tries to determine the type of netmask expression used in the downloadable ACL. If it detects a wildcard netmask expression, it converts it to a standard netmask expression. This option is useful when you are uncertain how the RADIUS server is configured; however, wildcard netmask expressions with holes in them cannot be unambiguously detected and converted. For example, the wildcard netmask 0.0.255.0 permits anything in the third octet, but the device might not detect this expression as a wildcard netmask. • Wildcard—The security appliance assumes that all downloadable ACLs received from the RADIUS server contain only wildcard netmask expressions, which it converts to standard netmask expressions.

AAA Server Dialog Box—TACACS+ Settings

Use the TACACS+ settings in the AAA Server dialog box to configure a TACACS+ AAA server object.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **TACACS+** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-9 AAA Server Dialog Box—TACACS+ Settings

Element	Description
Key Confirm	The shared secret that is used to encrypt data between the client and the AAA server. The key is a case-sensitive, alphanumeric string of up to 127 characters (U.S. English). Spaces and special characters are permitted. The key you define in this field must match the key on the TACACS+ server. Enter the key again in the Confirm field. Note the following: • Activity validation fails if you try defining a key with a space on a PIX, ASA, or FWSM device. • If you do not define a key, all traffic between the AAA server and its AAA clients is sent unencrypted.
Server Port	The port used for communicating with the AAA server. The default is 49.

AAA Server Dialog Box—Kerberos Settings

Use the Kerberos settings in the AAA Server dialog box to configure a Kerberos AAA server object.

Note This type of AAA server can be configured only on ASA, PIX 7.x+, and FWSM 3.1+ devices.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **Kerberos** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-10 AAA Server Dialog Box—Kerberos Settings

Element	Description
Server Port	The port used for communicating with the AAA server. The default is 88.
Kerberos Realm Name	The name of the realm containing the Kerberos authentication server and ticket granting server (maximum of 64 characters, typically all uppercase). For example, EXAMPLE.COM.
Retry Interval	The interval between attempts to contact the AAA server. Values range from 1 to 10 seconds.

AAA Server Dialog Box—LDAP Settings

Use the LDAP settings in the AAA Server dialog box to configure an LDAP AAA server object.

Note This type of AAA server can be configured only on ASA, PIX 7.x+, FWSM 3.1+, and IOS devices.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **LDAP** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-11 AAA Server Dialog Box—LDAP Settings

Element	Description
Enable LDAP over SSL/Secure Communication	Whether to establish a secure SSL connection between the device and the LDAP server. **Tip** You must select this option when using a Microsoft Active Directory LDAP server in order to enable password management.
No Negotiation (IOS only.)	When selected, this checkbox precludes further negotiation and moves to accept the channels previously established and accepted.
Server Port	The port used for communicating with the AAA server. The default is 389.
Login Directory	The name of the username or directory object in the LDAP hierarchy used for authenticated binding (maximum of 128 characters). Authenticated binding is required by some LDAP servers (including the Microsoft Active Directory server) before other LDAP operations can be performed. This field describes the authentication characteristics of the device. These characteristics should correspond to those of a user with administrator privileges. This string is case-sensitive. Spaces are not permitted in the string, but other special characters are allowed. Typically, this is a username such as DOMAIN\Administrator. However, you can use the more traditional format too, for example, cn=Administrator,OU=Employees,DN=example,DN=com.
Login Password	The case-sensitive, alphanumeric password for accessing the LDAP server (maximum of 64 characters). Spaces are not allowed.
Encrypted (IOS)	Whether the login password is encrypted.

Table 6-11 AAA Server Dialog Box—LDAP Settings (continued)

Element	Description
LDAP Hierarchy Location	The base distinguished name (DN), which is the location in the LDAP hierarchy where the authentication server should being searching when it receives an authorization request. For example, OU=Cisco. The maximum length is 128 characters. The string is case-sensitive. Spaces are not permitted, but other special characters are allowed.
PIX/ASA/FWSM Tab	
LDAP Scope	The extent of the search the server should make in the LDAP hierarchy when it receives an authorization request. The available options are: • onelevel—Searches only one level beneath the base DN. This type of search scope is faster than a subtree search, because it is less comprehensive. This is the default. • subtree—Searches all levels beneath the base DN (that is, searches the entire subtree hierarchy). This option takes more time.
LDAP Distinguished Name	The Relative Distinguished Name attribute (or attributes) that uniquely identifies an entry on the LDAP server. Common naming attributes are Common Name (CN), sAMAccountName, userPrincipalName, and User ID (uid). The case-sensitive, alphanumeric string can be up to 128 characters. Spaces are not permitted in the string, but other special characters are allowed.
SASL MD5 Authentication SASL Kerberos Authentication Kerberos Server Group	These options establish a Simple Authentication and Security Layer (SASL) mechanism to authenticate an LDAP client (the ASA/PIX/FWSM device) with an LDAP server. If you do not select one of these options, the simple mechanism is used, and usernames and passwords are transmitted in clear text. You can define one or both SASL authentication mechanisms. When negotiating SASL authentication, the ASA/PIX/FWSM device retrieves the list of SASL mechanisms configured on the LDAP server and selects the strongest mechanism configured on both devices. • SASL MD5 Authentication—Whether to have the device send the LDAP server an MD5 value computed from the username and password. You must configure the LDAP server to store the user passwords in reversible manner, or the LDAP server will not be able to validate the passwords. • SASL Kerberos Authentication—Whether to have the device send the LDAP server the username and realm using the GSSAPI (Generic Security Services Application Programming Interface) Kerberos mechanism. This mechanism is stronger than the MD5 mechanism. If you select Kerberos, you must also enter the name of the Kerberos AAA server group used for SASL authentication. The maximum length is 16 characters.

Table 6-11 AAA Server Dialog Box—LDAP Settings (continued)

Element	Description
LDAP Server Type	The type of LDAP server used for AAA: • Auto-Detect—The ASA/PIX/FWSM device tries to determine the server type automatically. This is the default. • Microsoft—The LDAP server is a Microsoft Active Directory server. **Note** You must configure LDAP over SSL to enable password management with Microsoft Active Directory. • Sun—The LDAP server is a Sun Microsystems JAVA System Directory Server. • OpenLDAP—The server is an Open LDAP server. You can use this only with ASA/PIX 8.0+ devices. • Novell—The server is a Novell LDAP server. You can use this only with ASA/PIX 8.0+ devices.
LDAP Attribute Map	The LDAP attribute configuration to bind to the LDAP server. Enter the name of an LDAP attribute map policy object or click **Select** to select it from a list or to create a new object. LDAP attribute maps take the attribute names that you define and map them to Cisco-defined attributes. For more information, see Add and Edit LDAP Attribute Map Dialog Boxes, page 6-46.
Group Base DN	(Microsoft LDAP AD servers only.) The base designated name (DN) under which all user groups are defined. When the ASA contacts the AD server for user group membership, the search starts at this DN. All groups must reside under this DN in the LDAP directory hierarchy and no group can reside outside of this path, or the group will not be found. Specifying this location can decrease the time required to complete user group searches. The alphanumeric string is case-sensitive and can be up to 128 characters. Spaces are not permitted in the string, but other special characters are allowed. For example: DN=cisco,DN=com **Tip** If you do not specify the group base DN, the LDAP Distinguished Name setting is used as the starting point for group searches.
Group Search Timeout	(Microsoft LDAP AD servers only.) The maximum time to wait for a response from an Active Directory server queried for user group information, in seconds. The default is 10 seconds, the range is 1 to 300 seconds.
IOS Tab	
Secure Cipher	The encryption method to be used.
Attribute Map (IOS)	The name of the IOS attribute map the server employs.
Secure Trust Point	The name of a trust point for certificates.

Table 6-11 AAA Server Dialog Box—LDAP Settings (continued)

Element	Description
Authentication bind-first	You can configure the sequence of search and bind of an authentication request with this option. The default is search first and then bind.
No Authorization Required	No authorization required for authentication requests.
Authentication Compare	Select this checkbox to replace the bind request with compare request for authentication. By default authentication request is performed with bind request.
User Object Filter	Specify the search filter user attribute type to be used in a search request. This helps in filtering out the requested user being searched.

AAA Server Dialog Box—NT Settings

Use the NT settings in the AAA Server dialog box to configure an NT AAA server object.

Note This type of AAA server can be configured only on ASA, PIX 7.x+, and FWSM 3.1+ devices.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **NT** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-12 AAA Server Dialog Box—NT Settings

Element	Description
Server Port	The port used for communicating with the AAA server. The default is 139.
NT Authentication Host	The name of the authentication domain controller hostname (maximum of 16 characters).

AAA Server Dialog Box—SDI Settings

Use the SDI settings in the AAA Server dialog box to configure an SDI AAA server object.

Note This type of AAA server can be configured only on ASA, PIX 7.x+, and FWSM 3.1+ devices.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **SDI** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-13 AAA Server Dialog Box—SDI Settings

Element	Description
Server Port	The port used for communicating with the AAA server. The default is 5500.
Retry Interval	The interval between attempts to contact the AAA server. Values range from 1 to 10 seconds. The default is 10 seconds.
SDI Server Version	The SDI server version: • SDI-pre-5—All SDI versions before version 5.0 • SDI-5—SDI version 5.0 or later.
SDI pre-5 Secondary Server	(Optional) A secondary server to be used for authentication if the primary server fails when using an SDI version prior to 5.0. Enter the IP address or the name of a network/host object, or click **Select** to select an object or create a new one.

AAA Server Dialog Box—HTTP-FORM Settings

Use the HTTP-FORM settings in the AAA Server dialog box to configure an HTTP-Form AAA server object for single sign-on authentication (SSO).

Note This type of AAA server can be configured only on ASA, PIX 7.x+, and FWSM 3.1+ devices.

Navigation Path

Go to the Add or Edit AAA Server Dialog Box, page 6-33 and select **HTTP-FORM** in the Protocol field.

Related Topics

- Creating AAA Server Objects, page 6-32
- Understanding AAA Server and Server Group Objects, page 6-27
- AAA Server Group Dialog Box, page 6-49

Field Reference

Table 6-14 AAA Server Dialog Box—HTTP-Form Settings

Element	Description
Start URL	The URL from which the WebVPN server of the security appliance should retrieve an optional pre-login cookie. The maximum URL length is 1024 characters. The authenticating web server might execute a pre-login sequence by sending a Set-Cookie header along with the login page content. The URL in this field defines the location from which the cookie is retrieved. **Note** The actual login sequence starts after the pre-login cookie sequence.
Action URI	The Uniform Resource Identifier (URI) that defines the location and name of the authentication program on the web server to which the security appliance sends HTTP POST requests for single sign-on (SSO) authentication. The maximum length of the action URI is 2048 characters. **Tip** You can discover the action URI on the authenticating web server by connecting to the web server's login page directly with a browser. The URL of the login web page displayed in your browser is the action URI for the authenticating web server.
Username Parameter	The name of the username parameter included in HTTP POST requests for SSO authentication. The maximum length is 128 characters. At login, the user enters the actual name value, which is entered into the HTTP POST request and passed on to the authenticating web server.
Password Parameter	The name of the password parameter included in HTTP POST requests for SSO authentication. The maximum length is 128 characters. At login, the user enters the actual password value, which is entered into the HTTP POST request and passed on to the authenticating web server.
Hidden Values	The hidden parameters included in HTTP POST requests for SSO authentication. They are referred to as hidden parameters because, unlike the username and password, they are not visible to the user. The maximum length of the hidden parameters is 2048 characters. **Tip** You can discover the hidden parameters that the authenticating web server expects in POST requests by using an HTTP header analyzer on a form received from the web server.
Authentication Cookie Name	The name of the authentication cookie used for SSO by the security appliance. The maximum length is 128 characters. If SSO authentication succeeds, the authenticating web server passes this authentication cookie to the client browser. The client browser then authenticates to other web servers in the SSO domain by presenting this cookie.

Add and Edit LDAP Attribute Map Dialog Boxes

Use the Add and Edit LDAP (Lightweight Directory Access Protocol) Attribute Map dialog boxes to populate the attribute map with name mappings that translate Cisco LDAP attribute names to custom, user-defined attribute names.

If you are introducing a security appliance to an existing LDAP directory, your existing custom LDAP attribute names and values are probably different from the Cisco attribute names and values. Rather than renaming your existing attributes, you can create LDAP attribute maps that map your custom attribute names and values to Cisco attribute names and values. By using simple string substitution, the security appliance then presents you with only your own custom names and values. You can then bind these attribute maps to LDAP servers or remove them as needed. You can also delete entire attribute maps or remove individual name and value entries.

For more information regarding LDAP support on ASA, PIX, and FWSM devices, see Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28.

Navigation Path

Select **Manage > Policy Objects**, then select **LDAP Attribute Map** from the Object Type selector. Right-click inside the table and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Creating AAA Server Objects, page 6-32
- AAA Server Dialog Box—LDAP Settings, page 6-40

Field Reference

Table 6-15 Add and Edit LDAP Attribute Map Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Attribute Map table	The table shows the mapped values. Each entry shows the customer map name, Cisco map name, and the attribute mapping of customer name to Cisco name. • To add a mapping, click the **Add Row** button to open the Add and Edit LDAP Attribute Map Value Dialog Boxes, page 6-47. • To edit a mapping, select it and click the **Edit Row** button. • To delete a mapping, select it and click the **Delete Row** button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-15 Add and Edit LDAP Attribute Map Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit LDAP Attribute Map Value Dialog Boxes

Use the Add and Edit LDAP Attribute Map Value dialog boxes to populate the attribute map with value mappings that apply user-defined attribute values to the custom attribute name and to the matching Cisco attribute name and value.

Navigation Path

From the Add and Edit LDAP Attribute Map Dialog Boxes, page 6-46, click the **Add Row** button to add a new mapping, or select a row and click the **Edit Row** button.

Field Reference

Table 6-16 Add and Edit LDAP Attribute Map Value Dialog Boxes

Element	Description
Customer Map Name	The name of your attribute map that relates to the Cisco map.
Cisco Map Name	The Cisco attribute map name you want to map to the customer map name.
Customer to Cisco Map Value table	The mappings of customer names to Cisco names. • To add a mapping, click the **Add Row** button to open the Add and Edit Map Value Dialog Boxes, page 6-47. • To edit a mapping, select it and click the **Edit Row** button. • To delete a mapping, select it and click the **Delete Row** button.

Add and Edit Map Value Dialog Boxes

Use the Add and Edit Map Value dialog boxes to map a customer LDAP attribute value to a Cisco map value. Enter the value from your LDAP map that you want to equate with a Cisco value.

Navigation Path

From the Add and Edit LDAP Attribute Map Value Dialog Boxes, page 6-47, click the **Add Row** button to add a new mapping, or select a row and click the **Edit Row** button.

Creating AAA Server Group Objects

You can create AAA server group objects for Security Manager policies requiring AAA services, such as authentication and authorization. Each AAA server group object can contain multiple AAA servers, all of which use the same protocol, such as RADIUS or TACACS+. For example, if you want to use RADIUS to authenticate network access and TACACS+ to authenticate CLI access, you must create at least two AAA server group objects, one for RADIUS servers and one for TACACS+ servers.

In addition, only one source interface can be defined for the AAA servers in the group. An error is displayed when you submit your changes if different AAA servers in the group use different source interfaces.

Note The error is triggered by the actual interface defined as the source, not the name of the interface role that represents the interface. That is, two AAA servers can have different interface roles defined as the source interface as long as they both resolve to the same device interface. An error is also displayed if the interface role defined for the source interface matches more than one actual interface on the device.

The number of AAA server group objects that can be created and the number of AAA server objects that can be included in each group object depend on the selected platform. For example, ASA devices support up to 18 single-mode server groups (with up to 16 servers each) and 7 multi-mode server groups (with up to 4 servers each). PIX firewalls support up to 14 server groups, each containing up to 14 servers.

Note Security Manager includes a predefined AAA server group object that you can use when you perform authentication locally inside the Cisco IOS router.

Tip You can also create AAA server group objects when you define policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Related Topics

- Creating Policy Objects, page 6-9
- Predefined AAA Authentication Server Groups, page 6-30
- Default AAA Server Groups and IOS Devices, page 6-31
- Understanding AAA Server and Server Group Objects, page 6-27

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **AAA Server Groups** from the Object Type selector.

Step 3 Right-click inside the work area, then select **New Object** to open the AAA Server Group Dialog Box, page 6-49.

Step 4 Enter a name for the object. The maximum name length is 16 characters if you plan to use this object with ASA, PIX, or FWSM devices and 128 characters for Cisco IOS routers. Spaces are not supported.

Note Cisco IOS routers do not support the following AAA server group names: RADIUS, TACACS, TACACS+. In addition, we do not recommend using an abbreviation of one of these names, such as rad or tac.

Step 5 Select the protocol to be used by the servers in the group.

Step 6 Enter the names of the AAA server policy objects that define the AAA servers to include in the group. Click **Select** to select the objects from a list filtered by the protocol you selected. You can also create new AAA server objects from the selection list. Separate multiple objects with commas.

Step 7 Configure the additional options that you want:

- Make this Group the Default AAA Server Group—For IOS devices only, whether you are using this group as the default group. Use this option if you intend to have a single global server group for this protocol for all policies requiring AAA. For more information, see Default AAA Server Groups and IOS Devices, page 6-31.
- ASA 8.4(2+) devices—If you are creating a RADIUS group containing Active Directory agent servers, select **AD Agent Mode**. This option indicates that the servers in the group are not full-function RADIUS servers but instead provide AD agent functions for identity-aware firewall. Use this group in the Identity Options policy.
- ASA, PIX, FWSM devices—Select options for how to handle AAA servers that stop responding, and for how to send accounting messages. For more information, see AAA Server Group Dialog Box, page 6-49.

Step 8 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 9 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 10 Click **OK** to save the object.

AAA Server Group Dialog Box

Use the AAA Server Group dialog box to create, copy, and edit AAA server groups. When defining a policy that uses a AAA server for authentication, authorization, or accounting, you select the server by selecting the server group to which the server belongs.

Navigation Path

Select **Manage > Policy Objects**, then select **AAA Server Groups** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Creating AAA Server Group Objects, page 6-48
- Understanding AAA Server and Server Group Objects, page 6-27
- Creating Policy Objects, page 6-9
- Add or Edit AAA Server Dialog Box, page 6-33
- Policy Object Manager, page 6-4

Field Reference

Table 6-17 AAA Server Group Dialog Box

Element	Description
Name	The object name (up to 16 characters when using this object with firewall devices; up to 128 characters for Cisco IOS routers). Object names are not case-sensitive. Spaces are not supported. Consider the following important points: • Cisco IOS routers do not support AAA server groups named RADIUS, TACACS, or TACACS+. In addition, we do not recommend using an abbreviation of one of these names, such as rad or tac. • If you define this AAA server group as the RADIUS or TACACS+ default group, any name you define here is automatically replaced in the device configuration by the default name (RADIUS or TACACS+) upon deployment.
Description	An optional description of the object.
Protocol	The protocol used by the AAA servers in the group. For more information about these options, see Supported AAA Server Types, page 6-28 and Additional AAA Support on ASA, PIX, and FWSM Devices, page 6-28.
AAA Servers	The AAA server policy objects that comprise the server group. Enter the names of the objects or click **Select** to select them from a list that is filtered to show only those AAA server objects that use the selected protocol. Separate multiple objects with commas. You can also create new objects from the selection list.
Make this Group the Default AAA Server Group (IOS) (IOS devices only.)	Whether to designate this AAA server group as the default group for the RADIUS or TACACS+ protocol. Select this option if you intend to use a single global group for the selected protocol for all policies on a specific device requiring AAA. Do not select this option if you intend to create multiple RADIUS or TACACS+ AAA server groups. Multiple groups can be used to separate different AAA functions (for example, use one group for authentication and a different group for authorization) or to separate different customers in a VRF environment. **Note** When you discover an IOS router, any AAA servers in the device configuration that are not members of a AAA server group are placed in special groups called CSM-rad-grp (for RADIUS) and CSM-tac-grp (for TACACS+), both of which are marked as default groups. These two groups are created solely to enable Security Manager to manage these servers. During deployment, the AAA servers in these special groups are deployed back to the device as individual servers. For more information, see Default AAA Server Groups and IOS Devices, page 6-31.

Table 6-17 AAA Server Group Dialog Box (continued)

Element	Description
AD Agent Mode (ASA 8.4(2+) devices only.)	Whether the servers in the group are Active Directory agents, which are used in identity-aware firewall configurations. You must select this option for an AD agent group to indicate that the group is not a full-function RADIUS server group. Use the AD agent group in the Identity Options policy. For more information, see Identifying Active Directory Servers and Agents, page 13-8.
Dynamic Authorization (ASA 9.2(1+) devices only.)	When using the RADIUS protocol, select the Dynamic Authorization check box to enable the RADIUS Dynamic Authorization Change of Authorization (CoA) services for the AAA server group. Specify the listening port for RADIUS CoA requests in the Port field. The valid range is 1024 to 65535 and the default value is 1700. Once defined, the corresponding RADIUS server group will be registered for CoA notification and the ASA will listen to the port for the CoA policy updates from the Cisco Identity Services Engine (ISE).
Interim Account Update (ASA 9.2(1+) devices only.)	When using the RADIUS protocol, select the Interim Account Update check box to enable the generation of RADIUS interim-accounting-update messages. Currently these messages are only generated when a VPN tunnel connection is added to a clientless VPN session. When this happens the accounting update is generated in order to inform the RADIUS server of the newly assigned IP address. Specify the length, in hours, of the interval between periodic accounting updates in the Interval field. The valid range is 1 to 120 and the default value is 24.
Authorize only (ASA 9.2(1+) devices only.)	When using the RADIUS protocol, select the Authorize only check box to enables authorize-only mode for the RADIUS server group. When this check box is selected, the common password configured for individual AAA servers is not required and does not need to be configured.
Max Failed Attempts (PIX, ASA, FWSM devices only.)	The number of connection failures that will be tolerated for any given server in the server group before that server is deactivated. The default is 3 attempts, the range is 1 to 5.
Internal Realm ID (ASA 9.8(1) and above devices only)	Enter a realm ID that corresponds to the RADIUS or LDAP protocol for the AAA server group policy object. **Note** The realm ID is a unique value in the range of 1-65535; it is only applicable for RADIUS and LDAP protocols.

Table 6-17 AAA Server Group Dialog Box (continued)

Element	Description
Reactivation Mode (PIX, ASA, FWSM devices only.)	The method to use when reactivating failed servers in the group: • Depletion—Reactivate failed servers only after all of the servers in the group are inactive. This is the default. When a server is deactivated, it remains inactive until all other servers in the group are inactive. When and if this occurs, all servers in the group are reactivated. This approach minimizes the occurrence of connection delays due to failed servers. If you configured a fallback method using the local database (for management access only) and all the servers in the group fail to respond, then the group is considered to be unresponsive, and the fallback method is tried. You can configure the Reactivation Deadtime value to determine the number of minutes that will elapse between the disabling of the last server in the group and the subsequent re-enabling of all servers. If you do not have a fallback method, the device continues to retry the servers in the group. • Timed—Reactivate failed servers after 30 seconds of downtime. This option is useful if the first server in the group is the primary server and you prefer that it be used whenever possible rather than the backup servers. This policy breaks down in the case of UDP servers. Because a connection to a UDP server will not fail, even if the server is not present, UDP servers are put back on line blindly. This could lead to slowed connection times or connection failures if a server group contains multiple servers that are not reachable.
Reactivation Deadtime (PIX, ASA, FWSM devices only.)	When you select **Depletion** as the reactivation mode, the number of minutes that should elapse between the deactivation of the last server in the group and the reactivation of all the servers in the group. The default is 10, the range is 0 to 1440 minutes (24 hours).
Group Accounting Mode (PIX, ASA, FWSM devices only.)	When using the RADIUS or TACACS+ protocols, the method for sending accounting messages to the AAA servers in the group: When using the server group for accounting (the protocol must be RADIUS or TACACS+), the method for sending accounting messages to the AAA servers in the group: • Single—Accounting messages are sent to a single server in the group. This is the default. • Simultaneous—Accounting messages are sent to all servers in the group simultaneously. If you select this option, the ASA forces the use of **Timed** as the reactivation mode.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-17 AAA Server Group Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Creating Access Control List Objects

An Access Control List (ACL) object is made up of one or more access control entries (ACEs), one or more ACL objects, or a combination of both. Each ACE is an individual permit or deny statement within an ACL. You can use ACL policy objects in several other policies and policy objects.

Beginning with Cisco Security Manager version 4.13, if an object group contains VM attribute and if it is applied to any other policies (except Access Rules), the deployment will fail. The VM Attribute Object is only applicable to the devices, which are ASA 9.7.1 or higher, and when object-group-search access-control is enabled.

You can create the following types of ACL objects:

- Extended – Extended ACLs enable you to specify source and destination addresses and service (or traffic protocol), and, based on the protocol type, the ports (for TCP or UDP), or the ICMP type (for ICMP) can be specified. For information on extended ACL objects, see Creating Extended Access Control List Objects, page 6-54.
- Standard – Standard ACLs use the source address for matching traffic. For information on standard ACL objects, see Creating Standard Access Control List Objects, page 6-56.
- Web – Web ACLs use destination address and port or a URL filter. For information on Web Type ACL objects, Creating Web Access Control List Objects, page 6-57.
- Unified – Unified ACL objects let you use source networks/hosts, source security groups, users, destination source networks/hosts, destination security groups, and services to match traffic. Further, the network/host specifications can contain IPv4 addresses, IPv6 addresses, or a combination of both. (With the release of Security Manager 4.4 and the ASA 9.0+, the separate IPv4 and IPv6 addressing/objects were “unified.”) See Creating Unified Access Control List Objects, page 6-58 for more information these ACLs.
- Ethertype – EtherType ACLs apply to non-IP layer-2 traffic on bridge group member interfaces only, in routed and transparent modes. You can use these rules to permit or drop traffic based on the EtherType value in the layer-2 packet.With EtherType ACLs, you can control the flow of non-IP traffic across the device. See Configuring Transparent Firewall Rules, page 23-1

For reference information about the dialog boxes used with these objects, see Add or Edit Access List Dialog Boxes, page 6-59.

Creating Extended Access Control List Objects

Extended access control lists allow you to permit or deny traffic from specific IP addresses to specific destination IP address and port, and specify the protocol of the traffic, such as ICMP, TCP, UDP, and so forth. Extended ACLs range from 100 to 199, and for devices running Cisco IOS Software Release 12.0.1 and higher, 2000 to 2699.

Extended ACL example:

```
access-list 110 - Applied to traffic leaving the office (outgoing)
access-list 110 permit tcp 10.128.2.0 0.0.0.255 any eq 80
```

ACL 110 permits traffic originating from any address on the 10.128.2.0 network. The "All-IPv4-Addresses" statement means that the traffic is allowed to have any destination address with the limitation of going to port 80. The value of 0.0.0.0/255.255.255.255 can be specified as "All-IPv4-Addresses."

Uses:

- Identifying addresses for NAT (policy NAT and NAT exemption)—Policy NAT lets you identify local traffic for address translation by specifying the source and destination addresses and ports in an extended access list. Regular NAT can only consider local addresses. An access list that is used with policy NAT cannot be configured to deny an access control entry (ACE).
- Identifying addresses for IOS dynamic NAT—For user-defined ACLs, the NAT plug-in generates its own ACL CLIs when deducing NAT traffic from VPN traffic.
- Filtering traffic that will be intercepted by Network Admission Control (NAC).
- Identifying traffic in a traffic class-map for modular policy—Access lists can be used to identify traffic in a class-map, which is used for features that support Modular Policy Framework such as TCP and general connection settings, inspection, IPS, and QoS. You can use one or more access lists to identify specific types of traffic.
- For transparent mode, enabling protocols that are blocked by a routed mode security appliance, including BGP, DHCP, and multicast streams. Because these protocols do not have sessions on the security appliance to allow return traffic, these protocols also require access lists on both interfaces.
- Establishing VPN access—You can use an extended access list in VPN commands to identify the traffic that should be tunneled on the device for an IPsec site-to-site tunnel or to identify the traffic that should be tunneled on the device for a VPN client. Use in conjunction with the policy objects and settings shown in Table 6-18 on page 6-54:

Table 6-18 Policy Objects and Settings

Policy Object	Device	Purpose
VPN Topology	Any	Selecting Protected Networks.
ASA User Group	ASA	Inbound Firewall Policy; Outbound Firewall Policy; Filter ACL.
Traffic Flow	ASA, PIX 7+	Service Policy Rules (MPC). The traffic flow BB (class-map) uses Extended ACL as one of its traffic match types.
User Group	• IOS • Catalyst 6500/7600 • PIX 6.3	For Easy VPN, Split Tunnel ACL and Firewall ACL (IOS devices only).

Related Topics

- Creating Access Control List Objects, page 6-53
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Creating Policy Objects, page 6-9
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

Step 1 Choose **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 From the Object Type selector, select **Access Control Lists**.

The Access Control List page appears. The Extended tab is displayed by default.

Step 3 Right-click inside the work area, then select **New Object**.

The Add Extended Access List dialog box appears (see Add or Edit Access List Dialog Boxes, page 6-59).

Step 4 Enter a name for the object and optionally a description of the object.

Note Make sure that the name of the ACL Object is unique and is not the same name as the Firewall Rules ACL defined in the Firewall ACL Setting. For more information, see Firewall ACL Setting Dialog Box, page 16-26.

Step 5 Right-click inside the table in the dialog box, then select **Add**.

The Add Extended Access Control Entry dialog box appears.

Step 6 Create the access control entry:

- If you choose **Access Control Entry** for Type, specify the characteristics of the traffic that you want to match and whether you are permitting or denying the traffic. Enter the source addresses whence the traffic originates, the destination addresses whither the traffic travels, and the services that define the characteristics of the traffic. Click **Advanced** to define logging options. For detailed information about the fields on the dialog box, see Add and Edit Extended Access Control Entry Dialog Boxes, page 6-61.
- If you choose **ACL Object**, select the object in the available objects list and click **>>** to add it to the list of selected objects.

Step 7 Click **OK** to save your changes.

The dialog box closes and you return to the Add Extended Access List page. The new entry is shown in the table. If necessary, select it and click the up or down buttons to position it at the desired location.

Step 8 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 9 Click **OK** to save the object.

Creating Standard Access Control List Objects

A standard access control list allows you to permit or deny traffic from specific IP addresses. The destination of the packet and the ports involved can be anything. Standard IP ACLs range from 1 to 99.

Standard ACL example:

```
access-list 10 permit 192.168.2.0 0.0.0.255
```

Uses:

- Identifying OSPF route redistribution.
- Filtering users of a community string using SNMP.
- Configuring VLAN ACLs for a Catalyst 6500/7600 device.

Related Topics

- Creating Access Control List Objects, page 6-53
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Creating Policy Objects, page 6-9
- Understanding Networks/Hosts Objects, page 6-80

Step 1 Choose **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 From the Object Type selector, select **Access Control Lists**.

The Access Control List page appears.

Step 3 Click the **Standard** tab.

Step 4 Right-click inside the work area, then select **New Object**.

The Add Standard Access List dialog box appears (see Add or Edit Access List Dialog Boxes, page 6-59).

Step 5 Enter a name for the object and optionally a description of the object.

Note Make sure that the name of the ACL Object is unique and is not the same name as the Firewall Rules ACL defined in the Firewall ACL Setting. For more information, see Firewall ACL Setting Dialog Box, page 16-26.

Step 6 Right-click inside the table, then select **Add**.

The Add Standard Access Control Entry dialog box appears.

Step 7 Create the access control entry:

- If you choose **Access Control Entry** for Type, specify the characteristics of the traffic that you want to match and whether you are permitting or denying the traffic. Enter the source addresses whence the traffic originates and select logging options. For detailed information about the fields on the dialog box, see Add and Edit Standard Access Control Entry Dialog Boxes, page 6-64.
- If you choose **ACL Object**, select the object in the available objects list and click **>>** to add it to the list of selected objects.

Step 8 Click **OK** to save your changes.

The dialog box closes and you return to the Add Standard Access List dialog box. The new entry is shown in the table. If necessary, select it and click the up or down buttons to position it at the desired location.

Step 9 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 10 Click **OK** to save the object.

Creating Web Access Control List Objects

Web ACLs, also referred to as WebVPN, let you establish a secure, remote-access VPN tunnel to the security appliance using a web browser. There is no need for either a software or hardware client. WebVPN provides easy access to a broad range of web resources and both web-enabled and legacy applications from almost any computer that can reach HTTPS Internet sites. WebVPN uses Secure Socket Layer Protocol and its successor, Transport Layer Security (SSL/TLS) to provide a secure connection between remote users and specific, supported internal resources that you configure at a central site.

The following table presents examples of Web VPN ACLs.

Table 6-19 Examples of Web VPN ACLs

Action	Filter	Effect
Deny	url http://*.yahoo.com/	Denies access to all of Yahoo!
Deny	url cifs://fileserver/share/directory	Denies access to all files in the specified location.
Deny	url https://www.company.com/directory/file.html	Denies access to the specified file.
Permit	url https://www.company.com/directory	Permits access to the specified location
Deny	url http://*:8080/	Denies HTTPS access to anywhere via port 8080.
Deny	url http://10.10.10.10	Denies HTTP access to 10.10.10.10.
Permit	url any	Permits access to any URL. Usually used after an ACL that denies url access.

Uses:

- As a filter ACL in an ASA User Group policy object (under SSL VPN > Clientless).

Related Topics

- Creating Access Control List Objects, page 6-53
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Creating Policy Objects, page 6-9

Step 1 Choose **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 From the Object Type selector, select **Access Control Lists**.

The Access Control List page appears.

Step 3 Click the **Web** tab.

Step 4 Right-click inside the work area and select **New Object**.

The Add WebType Access List dialog box appears (see Add or Edit Access List Dialog Boxes, page 6-59).

Step 5 Enter a name for the object and optionally a description of the object.

Note Make sure that the name of the ACL Object is unique and is not the same name as the Firewall Rules ACL defined in the Firewall ACL Setting. For more information, see Firewall ACL Setting Dialog Box, page 16-26.

Step 6 Right-click inside the access control entry table and choose **Add**.

The Add Web Access Control Entry dialog box appears.

Step 7 Create the access control entry:

- If you choose **Access Control Entry** for Type, specify the characteristics of the traffic that you want to match and whether you are permitting or denying the traffic. You can filter based on the network destination of the traffic (Network Filter) or the web address (URL Filter). For detailed information about the fields on the dialog box, see Add and Edit Web Access Control Entry Dialog Boxes, page 6-65.
- If you choose **ACL Object**, select the object in the available objects list and click **>>** to add it to the list of selected objects.

Step 8 Click **OK** to save your changes.

The dialog box closes and you return to the Add WebType Access List page. The new entry is shown in the table. If necessary, select it and click the up or down buttons to position it at the desired location.

Step 9 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 10 Click **OK** to save the object.

Creating Unified Access Control List Objects

A unified access control list allows you to permit or deny traffic from specific networks, hosts, security groups, and users, destined for specific networks, hosts and security groups. You also specify the service(s) involved.

Related Topics

- Creating Access Control List Objects, page 6-53
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Creating Policy Objects, page 6-9

- Understanding Networks/Hosts Objects, page 6-80

Step 1 Choose **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 From the Object Type selector, select **Access Control Lists**.

The Access Control List page appears.

Step 3 Click the **Unified** tab.

Step 4 Right-click inside the work area, then select **New Object**.

The Add Unified Access List dialog box appears (see Add or Edit Access List Dialog Boxes, page 6-59).

Step 5 Enter a name for the object and optionally a description of the object.

Note Make sure that the name of the ACL Object is unique and is not the same name as the Firewall Rules ACL defined in the Firewall ACL Setting. For more information, see Firewall ACL Setting Dialog Box, page 16-26.

Step 6 Right-click inside the table in the dialog box, then choose **Add**.

The Add Unified Access Control Entry dialog box appears.

Step 7 Create the access control entry:

- If you choose **Access Control Entry** for Type, specify the characteristics of the traffic that you want to match and whether you are permitting or denying the traffic. Enter the source addresses whence the traffic originates and select logging options. For detailed information about the fields on the dialog box, see Add and Edit Unified Access Control Entry Dialog Boxes, page 6-67.
- If you choose **ACL Object**, select the object in the available objects list and click **>>** to add it to the list of selected objects.

Step 8 Click **OK** to save your changes.

The dialog box closes and you return to the Add Unified Access List dialog box. The new entry is shown in the table. If necessary, select it and click the up or down buttons to position it at the desired location.

Step 9 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 10 Click **OK** to save the object.

Add or Edit Access List Dialog Boxes

Use the Add and Edit Access List dialog boxes to define access control entries (ACEs) for an ACL object. From this page, you can change the order of the ACEs and ACL objects within the table, add or edit ACEs and ACL objects, and delete ACEs and ACL objects.

The title of the dialog box indicates the type of ACL you are creating: Extended, Standard, or Web Type. The dialog boxes are essentially the same, the difference being the columns displayed in the ACE table.

Navigation Path

Select **Manage > Policy Objects**, then select **Access Control Lists** from the Object Type selector. Select the tab for the type of ACL object you want to create, and then right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Creating Access Control List Objects, page 6-53
- Creating Extended Access Control List Objects, page 6-54
- Creating Standard Access Control List Objects, page 6-56
- Creating Web Access Control List Objects, page 6-57
- Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

Field Reference

Table 6-20 Add and Edit Access List Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 6-20 Add and Edit Access List Dialog Boxes (continued)

Element	Description
Access Control Entry table	The access control entries (ACEs) and ACL objects that are part of the ACL. The table displays the name of the entry or object, description, options, services, and other attributes of the entry. In the Permit column, a green checkmark indicates that the entry permits traffic (typically, the traffic is considered a match for the service you are defining), whereas a red circle with a slash indicates that traffic is denied (typically, the traffic is considered to not match, and the service you are defining is not applied to the denied traffic). The source and, if applicable, destination addresses can be host IP addresses, network addresses, or network/host policy objects. • To add an ACE, click the Add button and fill in the dialog box for the type of ACL you are creating: – Add and Edit Extended Access Control Entry Dialog Boxes, page 6-61 – Add and Edit Standard Access Control Entry Dialog Boxes, page 6-64 – Add and Edit Web Access Control Entry Dialog Boxes, page 6-65 • To edit an ACE, select it and click the Edit button. • To delete an ACE, select it and click the Delete button. • To change the position of an entry, select it and click the Up/Down arrow buttons as required. Entries are evaluated top to bottom, so correct positioning is crucial for you to get the results you intend.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit Extended Access Control Entry Dialog Boxes

Use the Add or Edit Extended Access Control Entry dialog box to add an access control entry (ACE) or an ACL object to an Extended ACL object.

Navigation Path

From the Add or Edit Access List Dialog Boxes, page 6-59 for Extended ACL objects, click the **Add** button in the ACE table, or select a row and click the **Edit** button.

Related Topics

- Creating Extended Access Control List Objects, page 6-54

- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Filtering Items in Selectors, page 1-45

Field Reference

Table 6-21 Add and Edit Extended Access Control Entry Dialog Boxes

Element	Description
Type	The type of entry you are adding. The fields on the dialog box change based on your selection. • Access Control Entry—You want to define an ACE. • ACL Objects—You want to include an existing ACL object. You are presented with a list of available ACL objects. Select the objects you want to include and click the >> button to move them to the list of selected objects. You can remove an object by selecting it and clicking <<. You can also edit objects in the selected objects list.
Action	The action to take on traffic defined in the entry: • Permit—The service associated with this ACL is applied to this traffic. That is, the traffic is permitted to use the service. • Deny—The service associated with this ACL is not applied to this traffic. If there are multiple ACLs configured for a service, denied traffic is typically compared to the next ACL in the list; if it matches no permit entry in any ACL for the service, the service is not applied to the traffic. Whether the traffic is dropped from the network depends on the service.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-21 Add and Edit Extended Access Control Entry Dialog Boxes (continued)

Element	Description
Source Destination	The source or destination of the traffic. You can enter more than one value by separating the items with commas. You can enter any combination of the following address types. For more information, see Specifying IP Addresses During Policy Definition, page 6-87. • Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list. (ASA 8.4(2+) only.) You can select FQDN network/host objects to select traffic based on fully-qualified host names. • Host IP address, for example, 10.10.10.100. • Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0. • A range of IP addresses, for example, 10.10.10.100-10.10.10.200. • An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).
Users	(ASA 8.4(2+) only.) The Active Directory (AD) usernames, user groups, or identity user group objects for the rule, if any. The user specification is conjoined to the source address to limit the match to user addresses within the source address range. You can enter more than one value by separating the items with commas. You can enter any combination of the following values. • Individual user names: NetBIOS_DOMAIN\username • User groups (note the double \): NetBIOS_DOMAIN\\user_group • Identity user group object names. Click **Select** to select objects, users, or user groups from a list or to create new objects. For more information, see: • Selecting Identity Users in Policies, page 13-21 • Configuring Identity-Based Firewall Rules, page 13-21 • Creating Identity User Group Objects, page 13-19
Services	The services that define the type of traffic to act on. You can enter more than one value by separating the items with commas. You can enter any combination of service objects and service types (which are typically a protocol and port combination). If you type in a service, you are prompted as you type with valid values. You can select a value from the list and press Enter or Tab. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.

Table 6-21 Add and Edit Extended Access Control Entry Dialog Boxes (continued)

Element	Description
Description	An optional description of the object.
Advanced button	Click this button to define logging options for the entry: • For PIX, ASA, and FWSM devices, you can enable: – Default logging—If a packet is denied, message 106023 is generated. If a packet is permitted, no message is generated. – Per ACE logging—If a packet is denied, message 106100 is generated. You can select the logging severity level for the messages, and the interval (in seconds from 1 to 600) for generating messages. • For IOS devices, when you enable logging, informational messages about packets that match the entry are sent to the console. You can also elect to include the input interface and source MAC address or VC in the logging output.

Add and Edit Standard Access Control Entry Dialog Boxes

Use the Add or Edit Standard Access Control Entry dialog box to add an access control entry (ACE) or an ACL object to a Standard ACL object.

Navigation Path

From the Add or Edit Access List Dialog Boxes, page 6-59 for Standard ACL objects, click the **Add** button in the ACE table, or select a row and click the **Edit** button.

Related Topics

- Creating Standard Access Control List Objects, page 6-56
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Filtering Items in Selectors, page 1-45

Field Reference

Table 6-22 Add and Edit Standard Access Control Entry Dialog Boxes

Element	Description
Type	The type of entry you are adding. The fields on the dialog box change based on your selection. • Access Control Entry—You want to define an ACE. • ACL Objects—You want to include an existing ACL object. You are presented with a list of available ACL objects. Select the objects you want to include and click the >> button to move them to the list of selected objects. You can remove an object by selecting it and clicking <<. You can also edit objects in the selected objects list.

Table 6-22 Add and Edit Standard Access Control Entry Dialog Boxes (continued)

Element	Description
Action	The action to take on traffic defined in the entry: • Permit—The service associated with this ACL is applied to this traffic. That is, the traffic is permitted to use the service. • Deny—The service associated with this ACL is not applied to this traffic. If there are multiple ACLs configured for a service, denied traffic is typically compared to the next ACL in the list; if it matches no permit entry in any ACL for the service, the service is not applied to the traffic. Whether the traffic is dropped from the network depends on the service.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Source	The source of the traffic. You can enter more than one value by separating the items with commas. You can enter any combination of the following address types. For more information, see Specifying IP Addresses During Policy Definition, page 6-87. • Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list. • Host IP address, for example, 10.10.10.100. • Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0. • A range of IP addresses, for example, 10.10.10.100-10.10.10.200. • An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).
Description	An optional description of the object.
Log Option	Whether to create log entries when traffic meets the entry criteria. ACL logging generates syslog message 106023 for denied packets. Deny packets must be present to log denied packets.

Add and Edit Web Access Control Entry Dialog Boxes

Use the Add or Edit Web Access Control Entry dialog box to add an access control entry (ACE) or an ACL object to a Web Type ACL object.

Navigation Path

From the Add or Edit Access List Dialog Boxes, page 6-59 for Web Type ACL objects, click the **Add** button in the ACE table, or select a row and click the **Edit** button.

Related Topics

- Creating Web Access Control List Objects, page 6-57
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5

- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Filtering Items in Selectors, page 1-45

Field Reference

Table 6-23 Add and Edit Web Access Control Entry Dialog Boxes

Element	Description
Type	The type of entry you are adding. The fields on the dialog box change based on your selection. • Access Control Entry—You want to define an ACE. • ACL Objects—You want to include an existing ACL object. You are presented with a list of available ACL objects. Select the objects you want to include and click the >> button to move them to the list of selected objects. You can remove an object by selecting it and clicking <<. You can also edit objects in the selected objects list.
Action	The action to take on traffic defined in the entry: • Permit—The service associated with this ACL is applied to this traffic. That is, the traffic is permitted to use the service. • Deny—The service associated with this ACL is not applied to this traffic. If there are multiple ACLs configured for a service, denied traffic is typically compared to the next ACL in the list; if it matches no permit entry in any ACL for the service, the service is not applied to the traffic. Whether the traffic is dropped from the network depends on the service.
Filter Destination	Whether the entry specifies a network filter (host or network address) or a URL filter (web site address). Your selection changes the fields on the dialog box. The fields are described below.
Destination (Network Filter only.)	The destination of the traffic. You can enter more than one value by separating the items with commas. You can enter any combination of the following address types. For more information, see Specifying IP Addresses During Policy Definition, page 6-87. • Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list. • Host IP address, for example, 10.10.10.100. • Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0. • A range of IP addresses, for example, 10.10.10.100-10.10.10.200. • An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).

Table 6-23 Add and Edit Web Access Control Entry Dialog Boxes (continued)

Element	Description
Ports (Network Filter only.)	The port numbers or port list policy objects that define the port the traffic uses, if you want to use port identification. You can enter more than one value by separating the items with commas. You can enter any combination of the following types: • Port list object. Enter the name of the object or click **Select** to select it from a list. You can also create new port list objects from the selection list. • Port number, for example, 80. • A range of ports, for example, 80-90.
URL Filter (URL Filter only.)	The Universal Resource Locator (URL), or web address, of the traffic. You can use an asterisk as a match-all wildcard. For example, http://*.cisco.com matches all servers on the cisco.com network. You can specify any valid URL.
Logging	The type of logging to use for this entry: • Select Log Disabled to not create log entries. • Select Default to use the default settings on the device. • All other available options enable logging and identify the log level that will be used.
Logging Interval	The interval of time, in seconds, used to generate logging messages, from 1 to 600. The default is 300. You can modify this field only if you select a logging level in the Logging field.
Time Range	The time range policy object that defines the time range associated with the entry. The time range defines the access to the device and relies on the device's system clock. For more information, see Configuring Time Range Objects, page 6-71. Enter the name of the object or click **Select** to select it from a list. You can also create new time range objects from the selection list. **Note** Time range is not supported on FWSM 2.x or PIX 6.3 devices.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	An optional description of the object.

Add and Edit Unified Access Control Entry Dialog Boxes

Use the Add or Edit Unified Access Control Entry dialog box to add an access control entry (ACE) or an ACL object to a Unified ACL object.

Navigation Path

From the Add or Edit Access List Dialog Boxes, page 6-59 for Unified ACL objects, click the **Add** button in the ACE table, or select a row and click the **Edit** button.

Related Topics

- Creating Unified Access Control List Objects, page 6-58
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Filtering Items in Selectors, page 1-45

Field Reference

Table 6-24 Add and Edit Unified Access Control Entry Dialog Boxes

Element	Description
Type	The type of entry; the fields in the dialog box change based on your choice: • Access Control Entry—You want to define an ACE. • ACL Objects—You want to include one or more existing ACL objects. You are presented with a list of available ACL objects. Select the objects you want to include and click the **>>** button to move them to the list of selected objects. You can remove an object by selecting it and clicking **<<**. You can also edit an object in the selected objects list.
Action	The action to take on traffic defined in the entry: • Permit—The Services associated with the ACE are applied to this traffic. That is, the traffic defined by this entry is permitted to use the Services. • Deny—The Services associated with this ACE are not applied to this traffic. If there are multiple ACLs configured for a service, denied traffic is typically compared to the next ACE in the list; if it matches no permit entry in any ACL for the service, the service is not applied to the traffic. Whether the traffic is dropped from the network depends on the service.

Table 6-24 Add and Edit Unified Access Control Entry Dialog Boxes (continued)

Element	Description
Source	Provide traffic sources for this rule; can be networks and hosts. You can enter values or object names, or Select objects, for one or more of the following: • Networks/Hosts – You can specify a various network, host and interface definitions, either individually or as objects. If you Select an interface object as a source, the dialog box displays tabs to differentiate between hosts/networks and interfaces. Enter more than one value in any of these fields by separating the items with commas or ranges. The "All-Address" objects do not restrict the rule to specific hosts, networks, or interfaces. These addresses are IPv4 or IPv6 addresses for hosts or networks, network/host objects, interfaces, or interface roles. **Note** (ASA 8.4.2+ only) You can only specify a fully qualified domain name (FQDN) by providing an FQDN network/host object, or a group object that includes an FQDN object. You cannot directly type in an FQDN. See Understanding Networks/Hosts Objects, page 6-80, Specifying IP Addresses During Policy Definition, page 6-87 and Understanding Interface Role Objects, page 6-73 for additional information about these definitions. **Note** Enter the IPv6 addresses as comma separated values only. Preview configuration displays an error when IPv6 address is provided as a range. All Source, Source SG, and Users specifications area combined to limit traffic matches to only those flows that include all source definitions. For example, specified user traffic originating from within a specified source address range.
Source SG	(ASA 9.0+ only) Enter or Select the name or tag number for one or more source Security Groups for the ACE, if any. For more information about security groups, see: • Selecting Security Groups in Policies, page 14-16 • Configuring TrustSec-Based Firewall Rules, page 14-17 • Creating Security Group Objects, page 14-14

Table 6-24 Add and Edit Unified Access Control Entry Dialog Boxes (continued)

Element	Description
Users	(ASA 8.4.2+ only) Enter or Select the Active Directory (AD) user names, user groups, or identity user group objects for the ACE, if any. The user specification is conjoined to the source address to limit the match to user addresses within the source address range. You can enter more than one value by separating the items with commas. You can enter any combination of the following values: • Individual user names: NetBIOS_DOMAIN\username • User groups (note the double \): NetBIOS_DOMAIN\\user_group • Identity user group object names. For more information, see: • Selecting Identity Users in Policies, page 13-21 • Configuring Identity-Based Firewall Rules, page 13-21 • Creating Identity User Group Objects, page 13-19
Destination	The source or destination of the traffic. You can enter more than one value by separating the items with commas. **Note** Enter the IPv6 addresses as comma separated values only. Preview configuration displays an error when IPv6 address is provided as a range. Provide traffic destinations, and optionally destination security groups (ASA 9.0+ only), for this ACE. As with the source entries, you can enter values or object names, or Select objects, for one or more destinations.
Destination SG	(ASA 9.0+ only) Enter or Select the name or tag number for one or more source Security Groups for the ACE, if any. For more information about security groups, see: • Selecting Security Groups in Policies, page 14-16 • Configuring TrustSec-Based Firewall Rules, page 14-17 • Creating Security Group Objects, page 14-14
Service	The services that define the type of traffic to act on. You can enter more than one value by separating the items with commas. You can enter or Select any combination of service objects and service types (which are typically a protocol and port combination). If you type in a service, you are prompted as you type with valid values. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.

Table 6-24 Add and Edit Unified Access Control Entry Dialog Boxes (continued)

Element	Description
Advanced button	Click this button to open the Advanced dialog box and define logging options for the ACE. For PIX, ASA, and FWSM devices, you can enable: • Default logging—If a packet is denied, message 106023 is generated. If a packet is permitted, no message is generated. • Per ACE logging—If a packet is denied, message 106100 is generated. You can select the logging severity level for the messages, and the interval (in seconds from 1 to 600) for generating messages.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	An optional description of the object.

Configuring Time Range Objects

Use the Add or Edit Time Range dialog box to create, edit, or copy a time range object.

You can create time range objects for use when creating time-based ACLs and some firewall rules. While similar to extended ACLs in function, time-based ACLs allow for access control based on time considerations. The time range applies to specific rules, and makes those rules active for the specific time period defined in the range. For example, you can implement a rule for typical work hours to allow or prevent certain types of access.

You can also use time range objects when defining ASA user groups to restrict VPN access to specific times during the week. For more information, see ASA Group Policies SSL VPN Settings, page 34-25.

Time range objects can rely on the device's system clock, but they work best when using Network Time Protocol (NTP) synchronization.

Navigation Path

Select **Manage > Policy Objects**, then select **Time Ranges** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Field Reference

Table 6-25 Time Range Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object (up to 1024 characters).

Table 6-25 Time Range Dialog Box (continued)

Element	Description
Start Time End Time	The overall starting and ending time for the time range object: • Start Now—Defines the time of deployment as the start time. • Never End—Defines no end time for the range. • Start At, End At—Defines a specific start or end date and time. Click the calendar icon to display a tool for selecting the date. Enter the time in the Time field using the 24-hour clock format, HH:MM.
Recurring Ranges	Recurring time periods that happen within the overall start and end times, if any. For example, if you want to create a time range object that defines work hours, you could select Start Now and Never End for the overall range, and enter a recurring range of weekdays from 08:00 to 18:00 hours. • To add a range, click the **New Recurring Range** button and fill in the Recurring Ranges Dialog Box, page 6-72. • To edit a range, select it and click the **Edit Recurring Range** button. • To delete a range, select it and click the **Delete Recurring Range** button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Recurring Ranges Dialog Box

Use the Recurring Ranges dialog box to add or edit recurring time intervals that are defined as part of a time range object. You can define as many recurring ranges as required.

Navigation Path

Go to the Add or Edit Time Range dialog box and click the **New Recurring Range** button under Recurring Ranges, or select a range and click **Edit Recurring Range**. See Configuring Time Range Objects, page 6-71.

Field Reference

Table 6-26 Recurring Ranges Dialog Box

Element	Description
Specify days of the week and times during which this recurring range will be active	Defines a recurring range that is based on specific days and times of the week. You can select from: • Every day • Weekdays • Weekends • On these days of the week—Select the specific days to include in the range. Also select the starting and ending time during the day. The default is all day.
Specify a weekly interval during which this recurring range will be active	Defines a recurring range for every week. Select the starting and ending day and time. For example, you can start the weekly period on Sunday and end it on Thursday.

Understanding Interface Role Objects

Interface Role objects have the following uses:

- Specifying multiple interfaces— Interface role objects allow you to apply policies to specific interfaces on multiple devices without having to manually define the names of each interface. Because most devices follow a standard naming convention for their interfaces, you can define a naming pattern that describes a particular interface type and then assign a policy to all interfaces matching that pattern.
- Zones—You use interface role objects to define the zones in a zone-based firewall rules policy.

For example, you might define an interface role with a naming pattern of DMZ*. When you include this interface role in a policy, the policy is applied to all interfaces whose name begins with “DMZ” on the selected devices. As a result, you can, for example, assign a policy that enables anti-spoof checking on all DMZ interfaces to all relevant device interfaces with a single action. Interface roles can refer to any of the actual interfaces on the device, including physical interfaces, subinterfaces, and virtual interfaces, such as loopback interfaces.

Interface roles serve as an indirection entity between interfaces on the one hand and policies on the other. This enables you to apply policies to particular device interfaces based on the assigned role. Additionally, if you change the naming convention used for a particular interface type, you do not need to determine which policies are affected by the change. All you do is edit the interface role.

Interface roles are especially useful when you apply policies to new devices. As long as the devices you are adding share the same interface naming scheme as existing devices, the relevant policies can be extended to them without the need to make additional assignments.

Security Manager includes the following predefined interface roles:

- All-Interfaces—Includes every interface defined on a device.
- Internal—Includes only specific interfaces that are meant to be on the inside of a network. See the object definition for a list.

- External—Includes only specific interfaces that are meant to be on the outside of a network. See the object definition for a list.
- Self—Does not include any interfaces. The Self interface role is specific to zone-based firewall rules policies. The Self zone is the router itself. You can use it to identify traffic originating from the router, or traffic directed to the router. It does not include traffic passing through the router.

The following topics describe how to work with interface role objects:

- Creating Interface Role Objects, page 6-74
- Specifying Interfaces During Policy Definition, page 6-76
- Using Interface Roles When a Single Interface Specification is Allowed, page 6-77
- Handling Name Conflicts between Interfaces and Interface Roles, page 6-78
- Chapter 22, "Managing Traffic Zones"

Creating Interface Role Objects

You can create interface role objects that represent one or more interfaces on devices. You can then use these roles when you define policies that require interfaces or zones. When you create an interface role object, you must define the naming pattern of the device interfaces to include in the object. Interface roles can refer to any of the actual interfaces on the device, including physical interfaces, subinterfaces, and virtual interfaces.

Tip You can also create interface role objects when you define policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Related Topics

- Creating Policy Objects, page 6-9
- Specifying Interfaces During Policy Definition, page 6-76
- Understanding Interface Role Objects, page 6-73
- Using Interface Roles When a Single Interface Specification is Allowed, page 6-77
- Managing Object Overrides, page 6-17
- Chapter 22, "Managing Traffic Zones"

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **Interface Roles** from the Object Type selector.

Step 3 Right-click in the work area, then select **New Object**.

The Interface Role dialog box appears.

Step 4 Enter a name for the object and optionally a description of the object. Names can be up to 128 characters, descriptions up to 1024.

Step 5 Enter one or more naming patterns for the interface role object. The names are the complete or partial names of interfaces, subinterfaces, and other virtual interfaces. Separate multiple name patterns with commas.

You can use these wildcards to create name patterns that apply to multiple interfaces:

- Use a period (.) as a wildcard for a single character. To use a period as part of the pattern itself, enter a backslash (\) before the period.
- Use an asterisk (*) as a wildcard for one or more characters at the end of the interface pattern.

  For example, **DMZ*** would include all interfaces whose name begins with “DMZ”, while **DMZ.** would match interfaces such as DMZ1 and DMZ2, but would not match DMZ10.

If the pattern does not include a wildcard, it must match the exact name of the interface. For example, the pattern **FastEthernet** will not match FastEthernet0/1 unless you include an asterisk at the end of the pattern.

Step 6 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 7 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 8 Click **OK** to save the object.

Interface Role Dialog Box

Use the Interface Role dialog box to create, copy, or edit an interface role object. Interface Role objects have the following uses:

- Specifying multiple interfaces— Interface role objects allow you to apply policies to specific interfaces on multiple devices without having to manually define the names of each interface.
- Zones—You use interface role objects to define the zones in a zone-based firewall rules policy.

Navigation Path

Select **Manage > Policy Objects**, then select **Interface Roles** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Creating Policy Objects, page 6-9
- Creating Interface Role Objects, page 6-74
- Using Interface Roles When a Single Interface Specification is Allowed, page 6-77
- Specifying Interfaces During Policy Definition, page 6-76
- Understanding Interface Role Objects, page 6-73
- Chapter 22, “Managing Traffic Zones”
- Policy Object Manager, page 6-4

Field Reference

Table 6-27 Interface Role Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object. A maximum of 1024 characters is allowed.

Table 6-27 Interface Role Dialog Box (continued)

Element	Description
Interface Name Patterns	The names to include in this interface role. The names are the complete or partial names of interfaces, subinterfaces, and other virtual interfaces. Separate multiple name patterns with commas. **Note** For firewall devices, use the name assigned to the interface (for example, Inside, Outside, or DMZ) and not the hardware port identifier (for example, Ethernet0). You can use these wildcards to create name patterns that apply to multiple interfaces: • Use a period (.) as a wildcard for a single character. To use a period as part of the pattern itself, enter a backslash (\) before the period. • Use an asterisk (*) as a wildcard for one or more characters at the end of the interface pattern. For example, DMZ* would include all interfaces whose name begins with "DMZ", while DMZ. would match interfaces such as DMZ1 and DMZ2, but would not match DMZ10. If the pattern does not include a wildcard, it must match the exact name of the interface. For example, the pattern "FastEthernet" will not match FastEthernet0/1 unless you include an asterisk at the end of the pattern.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Specifying Interfaces During Policy Definition

When you configure policies that require you to identify an interface, you have several options for specifying the interface:

- Enter the name of the interface manually, for example, Ethernet0.

  To manually specify a subinterface as part of a policy definition, you must enter a backslash (\) before the period. For example, Ethernet0\.1.

  If you enter the period without the backslash, Security Manager treats the period as a wildcard for a single character. For example, if you want to define Ethernet1/1.0 as part of an access rule, you need to enter **Ethernet1/1\.0**. If you enter **Ethernet1/1.0** instead, the name matches interfaces named Ethernet1/1.0 and Ethernet1/1/0, because the period on its own is treated as a wildcard.

- Enter the name of an interface role manually. For more information about interface roles, see Understanding Interface Role Objects, page 6-73.

- Select an interface or interface role from a list. By clicking **Select** next to the Interfaces field, you are prompted with a list of valid interface names and interface roles. Subinterfaces appear with a backslash before the period in their names.

  By selecting from a list, you can ensure that your entry is valid. For more information, see Selecting Objects for Policies, page 6-2.

When a policy allows multiple interfaces, separate entries with commas.

In policies and object selectors, icons distinguish between interfaces and interface roles. If you create interface roles with the same name as interfaces, be careful to select exactly what you want. Table 6-28 explains the icons.

Table 6-28 Icons for Interfaces and Interface Roles

Type	Icon
Interface	
Interface role If you can edit the role, a pencil image overlays the icon.	
Global "interface" on ASA 8.3+ devices, used for rules created as global instead of interface-specific.	

Related Topics

- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Configuring Firewall Device Interfaces, page 46-3
- Understanding Interface Role Objects, page 6-73
- Creating Interface Role Objects, page 6-74
- Using Interface Roles When a Single Interface Specification is Allowed, page 6-77

Using Interface Roles When a Single Interface Specification is Allowed

Interface role objects can match a variable number of actual interfaces defined on a device depending on how you define the role. Thus, for a particular device, an interface role might match zero, one, or more than one interface. When you use an interface role in a policy, Security Manager converts the role to commands that configure all interfaces defined on the device that match the role.

Many policies, however, require that you specify a single interface name. If you use an interface role in a situation where the policy allows a single interface name, you should define the interface role so that it matches a single interface. If you use an interface role that matches two or more interfaces on the device, Security Manager selects the first interface on the device that matches the role, which might not be the interface you desire (or that will work properly).

Related Topics

- Specifying Interfaces During Policy Definition, page 6-76
- Understanding Interface Role Objects, page 6-73
- Creating Interface Role Objects, page 6-74

Handling Name Conflicts between Interfaces and Interface Roles

Under normal circumstances, you can configure an interface role that has the same name as an actual interface on the device. If you use object selectors when defining policies (see Selecting Objects for Policies, page 6-2), both the interface and the interface role are listed as available choices, enabling you to select either option. If you type in this common name when you define a policy, Security Manager automatically associates the interface role with the policy, not the interface.

However, a naming conflict can occur under the following circumstances:

1. You type the name of an interface when defining a policy.
2. You later create an interface role that has the same name.
3. You type this name again when defining a policy.
4. You click **Select** to display the object selector, or **Save** to save the policy, or in some cases, **OK** to update the policy.

When this sequence of events occurs, the Interface Name Conflict dialog box opens automatically so that you can select whether you want to specify the interface or the interface role. The dialog box lists only those names for which there are conflicts.

Related Topics

- Specifying Interfaces During Policy Definition, page 6-76
- Understanding Interface Role Objects, page 6-73

Understanding Map Objects

The objects in the Maps folder in the Policy Object Manager allow you to configure class, parameter, and policy maps for inspection rules, zone-based firewall rules, or IPS, QoS and connection rules policies. The types of maps you can use with these policies depends on the operating system running on the device as well as the specific version number, so typically it is best to configure the maps when you are configuring the policies.

Tip Devices enforce unique names for all configured maps. For example, you cannot use the same name for an FTP and DNS class map on the same device. If you select maps with the same name for a device, Security Manager automatically adds a numerical suffix to the duplicate names, for example, dnsmap_1.

The Maps folder contains the following folders. Subfolders organize the maps based on whether they are used for inspection or web content filtering.

- Class Maps—Layer 7 class maps used for identifying traffic that you want to act on.
- Parameter Maps—Parameter maps that configure settings used in zone-based firewall rules policies or other maps.
- Policy Maps—Layer 7 policy maps used for identifying the action to take on selected traffic.

Also included in the Maps folder are entries for TCP Map objects (a Layer 4 object), Regular Expression objects, and Regular Expression Group objects.

The following sections describe the different types of maps in more detail.

Class Maps

Class maps are subordinate to policy maps. You cannot specify a class map directly in a device policy. Instead, you create a policy map to incorporate the class map. The class map itself defines the match conditions for the traffic that you want to target in an inspection rule or zone-based firewall rule.

- ASA/PIX 7.2 and higher, and FWSM devices—You can create class maps for the inspection of DNS, FTP, HTTP, IM, and SIP traffic. You also have the option of defining the traffic match directly in the policy map object, but if you create separate class maps, you can reuse them in more than one policy map.
- IOS 12.4(6)T and higher devices—You can create class maps for the inspection of IM applications (AOL, ICQ, MSN Messenger, Windows Messenger, and Yahoo Messenger), P2P applications (eDonkey, FastTrack, Gnutella, Kazaa2), H.323, HTTP, IMAP, POP3, SIP, SMTP, Sun RPC. You can also create class maps for filtering web content using the Local, N2H2, Trend, and Websense objects.

  Unlike the class maps used for ASA/PIX/FWSM, you must create separate class maps and refer to them from the related policy maps. You can use these policy maps in zone-based firewall inspection or content filtering rules. For more information, see these topics:

  - Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
  - Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36

To create class maps, see these topics:

- Configuring Class Maps for Inspection Policies, page 17-28
- Configuring Class Maps for Zone-Based Firewall Policies, page 21-19

To create the regular expressions and regular expression groups that you can use in class, parameter, and policy maps, see these topics:

- Add/Edit Regular Expressions, page 17-108
- Configuring Regular Expression Groups, page 17-108

Parameter Maps

Parameter maps define settings that you can use in zone-based firewall inspection or content filtering rules, or in other policy map objects.

- Inspection—You can create Inspection Parameter maps for general zone-based firewall rule parameters, or Protocol Info Parameter maps for use with IM application inspection.
- Content Filtering—You can create the following parameter maps to define web content filtering: Local, N2H2, Trend, URL Filter, URLF Glob, Websense.

Policy Maps

You can configure policy maps to alter the default actions of inspection or to configure web content filtering in zone-based firewall settings policies. Policy maps typically apply to applications that require special handling, perhaps due to embedded IP address information or the fact that the traffic opens secondary channels on dynamically assigned ports.

The policy map identifies the action to take on traffic that matches the conditions identified in the map. For most policy maps, you can specify traffic match conditions by referring to a class map. However, some policy maps require that you specify the match criteria within the policy map.

You can configure these types of policy maps:

- Inspection Rules—When configuring inspection rules, you can use Security Manager to create policy map objects for the following applications: DCE/RPC, DNS, ESMTP, FTP, GTP, H.323, HTTP, IM, IP options, IPsec, NetBIOS, SIP, Skinny, and SNMP. for more information, see Configuring Protocols and Maps for Inspection, page 17-22.
- Zone-Based Firewall Inspection Rules—When configuring zone-based firewall inspection rules, you can use Security Manager to create policy map objects for the following applications: H.323, HTTP, IM (includes AOL, ICQ, MSN Messenger, Windows Messenger, and Yahoo Messenger), IMAP, P2P (includes eDonkey, FastTrack, Gnutella, Kazaa2), POP3, SIP, SMTP, Sun RPC. For more information, see Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16.
- Zone-Based Firewall Content Filtering Rules—When configuring zone-based firewall content filtering rules, you can use Security Manager to create Web Filter policy maps. You can also configure HTTP policy maps to inspect HTTP traffic. For more information, see Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36.
- IPS, QoS and Connection Rules—When configuring this service policy, which is specific to PIX 7.x+ and ASA devices, you can customize TCP inspection using a TCP map. For more information, see Configuring TCP Maps, page 58-22 and Chapter 58, "Configuring Service Policy Rules on Firewall Devices".

Understanding Networks/Hosts Objects

Networks/Hosts objects are logical collections of IP addresses that represent networks, hosts, or both.

Note As of Security Manager 4.4, there are no longer separate IPv4 and IPv6 Networks/Hosts objects—there is now a single, unified Networks/Hosts object, which may accept IPv4 addresses, IPv6 addresses, or both (in the case of group objects). However, group objects containing a mixture of IPv4 and IPv6 addresses can be assigned only to policies on ASA 9.0.1 and later devices. See Policy Object Changes in Security Manager 4.4, page 1-10 for more information.

When you create a Networks/Hosts object, you must choose the type of object, which defines and limits the type of addresses the object can contain:

- Group – You can include combinations of any of the following types of addresses:
  - Networks or subnets, specified by IPv4 addresses and subnet masks, or IPv6 prefixes and prefix lengths.
  - Ranges of IPv4 or IPv6 network addresses.
  - Individual hosts, specified by IPv4 or IPv6 addresses (but not a domain name).
  - Other network/host objects, selected from a list of existing Networks/Hosts objects, including fully qualified domain name (FQDN) objects.
- FQDN – (ASA 8.4(2+) only) This object can contain a single host's fully qualified domain name, such as myhost.cisco.com. The device uses DNS to periodically resolve the FQDN to its IP address.
- Host – This object can contain a single host IPv4 or IPv6 address, such as 10.100.10.10 or 2001:DB8::0DB8:800:200C:417A.
- Attribute – This object can contain one or more policy based VM attribute agents, which allow a user to define network objects to filter traffic according to attributes associated with one or more Virtual Machines (VMs) in an VMware ESXi environment managed by VMware vCenter. Each VM attribute agent communicates with a single vCenter server.

- Address Range – This object can contain a single range of IPv4 or IPv6 addresses; the start and end addresses must be different, with the start being lower than the end.
- Network – This object can contain a single IPv4 network address and subnet mask, such as 10.100.10.0/24, or a single IPv6 prefix and prefix length, such as 2001:DB8::/32.

Networks/Hosts group objects make it easier to manage scalable policies. By using the associative capabilities of Networks/Hosts objects, you can expand your policies along with your network. For example, when you make changes to the list of addresses contained in a Networks/Hosts object, the changes propagate to all other Networks/Hosts objects, and to policies that refer to that Networks/Hosts object.

The host, network, and address range objects have special uses when used in policies for an ASA 8.3+ device. On these devices, you can configure object NAT rules in the policy object itself. If you use the object on other types of device, this NAT configuration is ignored.

The following topics describe how to work with Networks/Hosts objects:

- Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81
- Creating Networks/Hosts Objects, page 6-82
- Using Unspecified Networks/Hosts Objects, page 6-86
- Specifying IP Addresses During Policy Definition, page 6-87
- VM Attribute Policies, page 6-89

Contiguous and Discontiguous Network Masks for IPv4 Addresses

A network mask determines which portion of an IPv4 address identifies the network and which portion identifies the host. Like the IP address, the mask is represented by four octets. (An octet is an 8-bit binary number equivalent to a decimal number in the range 0-255.) If a given bit of the mask is 1, the corresponding bit of the IP address is in the network portion of the address, and if a given bit of the mask is 0, the corresponding bit of the IP address is in the host portion.

Standard, or contiguous, network masks start with zero or more 1s followed by zero or more 0s. This kind of network mask is considered contiguous because it represents a network that consists of a contiguous IP address range. For example, the network 192.168.1.0/255.255.255.0 contains all the IP addresses ranging from 192.168.1.0 to 192.168.1.255.

The following table shows different methods of representing commonly used standard network masks:

Table 6-29 Standard Network Masks

Dotted Decimal Notation	Classless Inter-Domain Routing (CIDR) Notation
255.0.0.0	/8
255.255.0.0	/16
255.255.255.0	/24
255.255.255.255	/32

For example, 255.255.255.0 indicates that the first three octets of the IP address (24 bits or /24 in CIDR notation) are made up of ones and identify the network; the last octet is made up of zeros and identifies the host.

Discontiguous Network Masks

Nonstandard, or discontiguous, network masks are masks that do not conform to the contiguous format. For example, 10.0.1.1/255.0.255.255 indicates that you want to match an address that matches octets 1, 3, and 4 exactly, but any value in octet 2 is accepted.

Although discontiguous network masks are not typically used for network configurations, they are sometimes used for certain commands, such as filtering commands when defining access control lists (ACLs). Security Manager supports the use of nonstandard network masks in the policies whose CLI commands support them. An error is displayed if you try to define a discontiguous network mask in a policy that does not support them.

Network Masks and Discovery

During discovery, Security Manager attempts to match network/host objects with existing equivalent objects defined in the Policy Object Manager:

- For contiguous network masks—Two network/host objects containing only standard networks are considered equivalent if they consist of the same set of IP addresses.
- For discontiguous network masks—Two network/host objects are considered equivalent only if the standard networks consist of the same set of IP addresses and the nonstandard networks are syntactically equivalent.

How Network Masks are Displayed

Although you can enter both contiguous and discontiguous network masks using dotted decimal notation, all contiguous network masks are converted to CIDR notation. This makes it easier to distinguish them from discontiguous network masks, which are displayed in dotted decimal notation only.

Related Topics

- Creating Networks/Hosts Objects, page 6-82
- Specifying IP Addresses During Policy Definition, page 6-87
- Using Unspecified Networks/Hosts Objects, page 6-86
- Understanding Networks/Hosts Objects, page 6-80

Creating Networks/Hosts Objects

You can create Networks/Hosts objects to represent networks, individual hosts, or groups of both. When you create a Networks/Hosts object, you must choose the type of object (group, host, FQDN, network, attribute, address range). Once created, you cannot change the object type.

Tip You can create Networks/Hosts objects "on the fly" when defining policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

You can specify NAT object only if you have the Modify privilege mapped to your role.

Related Topics

- Understanding Networks/Hosts Objects, page 6-80
- Creating Policy Objects, page 6-9
- Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81
- Specifying IP Addresses During Policy Definition, page 6-87

- VM Attribute Policies, page 6-89
- Using Unspecified Networks/Hosts Objects, page 6-86
- How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups, page 6-107

Step 1 Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window; see Policy Object Manager, page 6-4 for more information.

Step 2 Select **Networks/Hosts** in the Object Type selector.

Step 3 Click the New Object button at the bottom of the window and choose one of the following types of Networks/Hosts object to open the Add or Edit Network/Host Dialog Box, page 6-83. You also can right-click in the work area, choose **New Object**, and then choose one of the following options to open the dialog box.

- **Group** – To create an object that has one or more entry. You can include any combination of networks, hosts, address ranges, or other network/host objects (including FQDN objects).
- **FQDN** – (ASA 8.4(2+) only) To create an object with a single host's fully qualified domain name, such as myhost.cisco.com.
- **Host** – To create an object with a single host address, such as 10.100.10.10 or 2001:DB8::12ab:5689.
- **Attribute** – (ASA 9.7.1+ only) To create a network object to filter traffic according to attributes associated with one or more virtual machines (VMs) in a VMware ESXi environment managed by VMware vCenter.
- **Address Range** – To create an object with a single range of addresses, such as 10.100.10.1-10.100.10.255.
- **Network** – To create an object with a single network address, such as 10.100.10.0/24 or 2001:DB8::/32.

Tip Host, network, and address range objects also let you configure object NAT rules for ASA 8.3+ devices. Any NAT configuration is ignored for other devices.

Step 4 Provide the appropriate information in the Add or Edit Network/Host Dialog Box, page 6-83.

Add or Edit Network/Host Dialog Box

Use the Add or Edit Network/Host dialog box to view, create, or edit network/host objects. The title, content and appearance of the dialog box differ slightly based on the type of network/host object you are creating: Group, FQDN, Host, Attribute, Address Range, or Network. FQDN objects require ASA 8.4.2 or later devices. Attribute objects require ASA 9.7.1 or later devices.The Group type lets you enter multiple definitions, so you can have a collection of networks, hosts, and other network/host objects, whereas the other types allow a single definition only.

The Host, Network, and Address Range versions of the dialog box provide two tabbed panels of options: General and NAT. Options on the General panel and the non-tabbed versions of the dialog box are described in the following table; the NAT options are described in Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42.

Note As of Security Manager 4.4, there are no longer separate IPv4 and IPv6 Networks/Hosts objects—there is now a single, unified Networks/Hosts object, which may accept IPv4 addresses, IPv6 addresses, or both (in the case of group objects only). However, group objects containing a mixture of IPv4 and IPv6 addresses can be assigned only to policies on ASA 9.0.1 and later devices.

When you create IPv4-based Host, Network, or Address Range objects for use on ASA 8.3+ devices, or unified Host, Network, or Address Range objects for use on ASA 9.0.1+ devices, you can also configure object NAT rules on the NAT tab of the dialog box. In both cases, you must select **Allow Value Override per Device** to allow object NAT. For reference information on the NAT tab, see Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42.

In addition, you can create an object with no addresses. For this type of object, you must also select **Allow Value Override per Device** and create overrides for every device that uses the object. For more information about using unspecified addresses, see Using Unspecified Networks/Hosts Objects, page 6-86.

Navigation Path

Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Networks/Hosts** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

Related Topics

- Creating Networks/Hosts Objects, page 6-82
- Understanding Networks/Hosts Objects, page 6-80
- Policy Object Manager, page 6-4
- How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups, page 6-107
- Filtering Items in Selectors, page 1-45

Field Reference

Table 6-30 Network/Host Dialog Box (General Tab)

Element	Description
Name	The object name (up to 64 characters). Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-30 Network/Host Dialog Box (General Tab) (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. **Tip** If you configure NAT for host, address range, or network objects, you must select this option. The NAT configuration is created as a device override and is not kept in the object. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.
Group object options	
Available Networks/Hosts Members In Group Type in comma separated IP addresses	The Members In Group list shows the networks, hosts, and other network/host objects that are included in this object. To populate the list, do any combination of the following: • Select one or more Address, Attribute, FQDN, Group, Host, Network objects in the Existing Networks/Hosts list and then click the >> button between the lists. • Type one or more IP addresses in the "Type in comma separated IP addresses" field and then click the >> button between the lists. Separate multiple addresses with commas; they are added as separate lines in the Members list. For IPv4 addresses, you can include host addresses, network addresses (with subnet masks entered after a / character, such as 10.100.10.0/24), or a range of addresses (separate the starting and ending address with a hyphen, and optionally include a subnet mask). For IPv6 addresses, you can include host addresses, network addresses (with prefixes entered after a / character, such as 2001:DB8::/32), or a range of addresses (such as 2001:DB8::1-2001:DB8::100). See Specifying IP Addresses During Policy Definition, page 6-87 for more information. • To remove an item from the Members In Group list, select it and click the appropriate << button to return the item to its source location. You can select and remove multiple items at one time. **Note** Group objects containing a mixture of IPv4 and IPv6 addresses can be assigned only to policies on ASA 9.0.1 and later devices.
FQDN object options	
FQDN FQDN Type	The fully qualified domain name of a single host; for example, somehost.cisco.com. The FQDN Type specifies the type of IP address mapped to the provided domain: IPv4 Only, IPv6 Only, or **Default**, which applies a device-specific default; for all non-ASA and pre-9.0.1 ASA devices, the default is IPv4.

Table 6-30 Network/Host Dialog Box (General Tab) (continued)

Element	Description
Host object options	
IP Address	The IPv4 or IPv6 address of the single host to include in the object.
Attribute object options	
Agent Name Type Value	The VM Attribute Agent name. Select from a list of VM attribute agents or add a new VM Attribute Agent. The VM Attribute Agent Type should not exceed 128 characters. The VM Attribute Agent Value should not exceed 128 characters. **Note** A user can assign custom attribute types and values to a set of VMs in order to apply a common set of policies to a set of VMs with a common user-defined characteristic
Address Range object options	
Start IP Address End IP Address	The first and last IP address that define a range of addresses. The start and end addresses must be different, with the start being lower than the end.
Network object options	
IP Address Net Mask/Prefix	The IPv4 or IPv6 address that represents the network; for example, 10.100.10.0 or 2001:DB8::/32. If you entered an IPv4 address, enter its subnet mask in the Net Mask/Prefix field. You can type a mask in either CIDR format, for example, 24 (without the forward slash), or in dotted decimal format, for example, 255.255.255.0. If you entered an IPv6 address, enter its prefix length in the Net Mask/Prefix field.

Using Unspecified Networks/Hosts Objects

When you define a Networks/Hosts object, you can leave the address fields blank, thereby creating a Networks/Hosts object with an unspecified value. Networks/Hosts objects with unspecified values require that you create overrides for every device that uses them.

The advantage of using a Networks/Hosts object with an unspecified value is that Security Manager displays an error if you submit your changes without creating a device-level override on every device using the object. By contrast, when you define the global object with a placeholder value (such as, 10.10.10.10), that global value could be deployed by mistake if you fail to define an override.

The following procedure describes how to create and implement Networks/Hosts objects with unspecified values.

Related Topics

- Creating Networks/Hosts Objects, page 6-82
- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81
- Specifying IP Addresses During Policy Definition, page 6-87

- Understanding Networks/Hosts Objects, page 6-80

Step 1 Create a Networks/Hosts object, making sure to:

- Leave the address fields blank (for example, the Members in Group, IP Address and Net Mask/Prefix, FQDN, or Start and End IP Address).
- Select the **Allow Value Override per Device** check box.

For more information, see Creating Networks/Hosts Objects, page 6-82.

Step 2 Create overrides for each device that will use the object:

a. Click the green checkmark in the Overrides column for the object in the Networks/Hosts table to open the Policy Object Overrides Window, page 6-20.

b. Click the **Create Override** button and select the devices on which you want to create overrides, then define a value in the address field. At this point, this override value applies to all the selected devices. For more information, see Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19.

c. Double-click each device in the Policy Object Overrides dialog box, then modify the address field for the value required by that device.

Step 3 Define a policy that requires this object. You can use one of two methods:

- Define the policy on a single device in Device view, share the policy, then assign the policy to the other devices. See Sharing a Local Policy, page 5-41 and Modifying Shared Policy Assignments in Device View or the Site-to-Site VPN Manager, page 5-49.
- Create a shared policy in Policy view, then assign the policy to the other devices using the Assignments tab. See Modifying Policy Assignments in Policy View, page 5-54.

Note You can create a Networks/Hosts group object that refers to a Networks/Hosts object with an unspecified value. You do not have to create the device-level overrides before you assign the policy containing the object to devices.

Specifying IP Addresses During Policy Definition

Many policies and policy objects require that you enter an IP address for a host or network. For some policies or objects, you must enter just a host, or just a network. For other policies or objects, you can enter some combination of hosts and networks. You are prevented from entering or selecting addresses that are not appropriate for the circumstances.

The following is a description of all acceptable formats that you can use, both for IPv4 and IPv6 addresses, although a particular policy or object might not allow specific formats (for example, interface roles are allowed as address designations in only a very limited number of policies). If the policy or object allows it, you can enter multiple addresses by separating them with commas.

- Networks/Hosts object. Enter the name of the object or click **Select** to select it from a list. You can also create new Networks/Hosts objects from the selection list.

Note The only way to specify a fully qualified domain name (FQDN) is to use an FQDN Networks/Hosts object, or a group object that includes an FQDN object. You cannot directly type in an FQDN.

- Host IP address, in v4 or v6 format.
  - Complete IPv4 address; for example, 10.10.10.100
  - Complete IPv6 address, showing all eight components. For example, 2001:DB8:0:0:0DB8:800:200C:417A. It is not necessary to include the leading zeros in an individual field. Security Manager converts the address to compressed format if possible.
  - Compressed IPv6 address, where a group of fields is replaced by two colons (::). It is common for IPv6 addresses to contain successive hexadecimal fields of zeros. To make IPv6 addresses less cumbersome, you can use two colons (::) to compress successive hexadecimal fields of zeros at the beginning, middle, or end of an IPv6 address (the colons represent successive hexadecimal fields of zeros). You can use :: at most once in an IPv6 address. For example, 2001:DB8::0DB8:800:200C:417A. The unspecified address, 0:0:0:0:0:0:0:0, can be represented as ::. The loopback address is ::1.
  - IPv6 representation of an IPv4 address. When dealing in mixed IPv4/IPv6 environments, you can represent the IPv4 addresses in an alternate IPv6 format: x:x:x:x:x:x:d.d.d.d, where the Xs are the hexadecimal values of the first 6 fields, and the Ds are the IPv4 address with the octets separated by periods. The first 6 fields are either all zeros, ::FFFF, or 2001:DB8::. For example, 0:0:0:0:0:0:10.1.68.3, which in compressed format is ::10.1.68.3, or 0:0:0:0:0:FFFF:10.1.68.3, or 2001:DB8::10.1.68.3.
- Network address, in either IPv4 or IPv6 format:
  - IPv4 address, including subnet mask, in either CIDR format (10.10.10.0/24), or dotted decimal format (10.10.10.0/255.255.255.0).
  - IPv6 address, including the prefix length in decimal format in a manner similar to CIDR notation for IPv4, for example, /64. The number specifies the number of the left-most contiguous bit of the address that comprise the prefix. For example, 2001:DB8:0:CD30::/60.

    Note You could also enter 2001:DB8:0:CD30::/60 as 2001::CD30:0:0:0:0/60. However, compressing the trailing zeros is the preferred method, and Security Manager will translate the address to 2001:DB8:0:CD30::/60.

    For more detailed information on IPv6 addressing, see the IETF RFC 4291, IP Version 6 Addressing Architecture, at http://www.ietf.org/rfc/rfc4291.txt.
- A range of IP addresses. Separate the beginning and ending addresses with a hyphen. The range does not need to be within a single subnet unless the policy requires it.

  You can also include a prefix or subnet mask in CIDR format; for example, 2001:db8::1 - 2001:db8::2/64, or 10.10.10.100-10.10.10.200/24.
- An IPv4 address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).
- Interface role object (in rare cases). Enter the name of the object or click **Select** to select it from a list (you must select Interface Role as the object type). When you use an interface role, the rule behaves as if you supplied the IP address of the selected interface. This is useful for interfaces that get their address through DHCP, because you do not know what IP address will be assigned to the device. For more information, see Understanding Interface Role Objects, page 6-73.

When you create a network/host object or define IP addresses as part of a policy, Security Manager verifies that the syntax of the address is correct and that a mask or prefix was entered when required. For example, when you define a policy that requires a host, you do not need to enter a mask/prefix. However, when you define a policy that requires a subnet, you must enter the address with the mask/prefix, or select a network/host object that has a mask/prefix defined.

Related Topics

- Creating Networks/Hosts Objects, page 6-82
- Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81
- Using Unspecified Networks/Hosts Objects, page 6-86
- Policy Object Manager, page 6-4
- Understanding Networks/Hosts Objects, page 6-80

VM Attribute Policies

You can define network objects to filter traffic based on attributes associated with one or more virtual machines (VMs) in a VMware ESXi environment. This environment is managed by VMware vCenter. Users can assign attributes to VMs within the ESXi environment and configure an attribute agent; the attribute agent connects to vCenter or to a single ESXi host using HTTPS and requests and retrieves one or more bindings that associate the specific attribute to the primary IP address of the an ESXi VM.

A single ASA can have multiple attribute agents defined; each communicating with a different vCenter, or one or more communicating with the same vCenter.

This enables a user to define access control lists (ACLs) to assign policies to traffic from a group of VMs that share one or more attributes.This feature is referred to as Policy Based on VM Attributes

The VM attributes feature is supported on all hardware platforms, and on all ASAv platforms running on ESXi, KVM, or HyperV hypervisors. VM attributes can only be retrieved from VMs running on an ESXi hypervisor.

Note The ASA uses the term **attribute** or **attribute type** to refer to the characteristic to be monitored. VMWare uses the term **property** for the same characteristic. The terms may be used interchangeably.

Communication between the VM attribute agent and vCenter

There are two types of messages exchanged between the VM attribute agent and the vCenter - Property Request and Binding Update.:

- **Property Request** -This is a HTTPS message sent from the ASA to the IP address of the vCenter Server, indicating the complete list of attribute types currently configured for network objects associated with this attribute agent. attributes that have been configured. This message contains the SSL credentials necessary to authenticate the connection to vCenter. The vCenter responds with a corresponding HTTPS response.
- **Binding Update** - This is a asynchronous HTTPS message sent from the vCenter to the ASA, whenever an attribute changes for one or more VMs. Each binding update is identified by the IP address of the VM reporting the attribute change. If multiple attributes are being monitored by a single agent, a single binding update contains the current value of all monitored attributes for each VM. If a specific attribute being monitored by the agent is not configured on a VM, the binding will contain an empty attribute value for that VM. If a VM has not been configured with any monitored attributes, vCenter does not send a binding update to the ASA.

When an attribute agent issues a property request containing a new attribute type, vCenter responds with a binding update for each VM where the attribute type is configured. After that point, vCenter only issues a new binding when an attribute value is added or changed on a VM.

Attribute Agent States

There are two kinds of attribute agent states - **Connection State** and **Agent State**.

- **Connection State** - This indicates whether or not the attribute agent is currently in contact with vCenter.

Table 6-31 Connection State Table

Connection State	Explanation
No Host Credentials	The user has not entered vCenter host credentials using the **host** subcommand, or the agent has been deleted using the **no attribute source-group** command while there are network objects still using the agent.
Disconnected	The agent has host credentials defined, but is currently not in contact with vCenter. The connection is established when the ASA receives a HTTP 200 response to a keepalive packet.
Connected	The agent has received a response from vCenter to the latest keepalive packet.
Invalid Host Credentials	The agent has attempted to contact vCenter to issue a property request, but the request was rejected because the user name and/or password was incorrect. The agent stays in this state until new credentials are entered, at which point it will move to Disconnected state until a keepalive response is received from vCenter.

- **Agent State** - This indicates whether or not any network objects are configured to monitor attribute types through this agent.

Table 6-32 Agent State Table

Agent State	Explanation
Inactive	The agent currently has no attributes configured.
Active	The agent has one or more attributes configured. An agent can be Active even if there is no connection to vCenter.

Guidelines for Configuring vCenter Virtual Machines

To leverage the VM attributes feature, those attributes must be made available to the vCenter server by the managed virtual machines. As an example, some attributes are:

- **summary.config.name** - The user-defined name associated with the virtual machine - for example, VM-build-machine-1

- **summary.config.guestFullName** - The full name of the guest OS running on the virtual machine - for example, Red Hat Enterprise Linux 7 (64-bit)
- **summary.config.annotation** - The text description field for the virtual machine.

For string attribute values such as **summary.config. annotation**, the value in the network object attribute definition must be an exact match to the value reported to vCenter by the VM. For example, a network object attribute value, 'This is a Build Machine' does not match the VM **summary.config.annotation** value, 'this is a build machine' on the VM. A binding update containing the latter string will not be added to the host-map for the former.

VMs being monitored by the VM attributes feature must have VMware Tools installed. VMware Tools is the software component that reports the IPv4 or IPv6 address of the VM to the vCenter server. Since the function of VM attribute is to bind an IP address to an attribute type/value pair, vCenter will not report any binding information for VMs that are not running VMware Tools.

Within the ESXi environment, VMs are defined by a primary IP address, which roughly corresponds to the management IP address of an ASA. There can only be one primary address per VM, which can be either an IPv4 or IPv6 address. Bindings are always provided between the primary IP address and the attribute type/value pair. If a VM is configured with multiple IP addresses (such as IPv6 link local addresses), vCenter will only send binding updates for the primary address (usually the first address configured).

Note A user can assign custom attribute types and values to a set of VMs in order to apply a common set of policies to a set of VMs with a common user-defined characteristic.

For a comprehensive list of the attributes and related guidelines, refer to the VMware vCenter 5.5/6.0 documentation.

Configuring VM Attribute Policies

There are three steps to configuring a policy based on VM attributes:

Step 1 Configure the Network Object Attribute.

a. Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Networks/Hosts** from the Object Type Selector. Right-click inside the work area and select **New Object > Attribute**; you also can use the **+** button at the bottom of the pane to add a new network object attribute.

Note A Network Attribute Object can be used only if object-group-search is enabled.

b. Select a **VM Attribute Agent**, specify a **VM Attribute Type** and add a value for **VM Attribute Value**.

Step 2 Add a VM Attribute Agent.

a. Specify a **Name** for the VM Attribute Agent and add a **Description** for the VM attribute Agent.

b. By default, the **Agent Type** is esxi.

c. Enter the primary IP address of the vCenter server in the **DNS Host Name/IP Address** field.

d. Specify a **Username** and **Password** to authenticate to the vCenter Server.

e. Specify a duration of time that the connection is kept active while the agent is contacting the vCenter server. The default value of the **Retry Interval** is 30 seconds.

f. Specify the number of times that the agent will attempt to contact the vCenter server before declaring it inactive in the **Retry Count** field. The default value is 3.

g. Click **OK**.

Step 3 Configure an access-list using a VM Attribute. For more information see Creating Access Control List Objects, page 6-53

Note VM Attribute only supports an access-list object.

Understanding Pool Objects

Pool objects have the following uses:

- Specifying pools for use in Layer 3 load balancing for ASA clusters
- Specifying pools for use in Layer 3 EIGRP and OSPFv3 on ASA clusters

The following topics describe how to work with pool objects:

- Add or Edit IPv4 Pool Dialog Box, page 6-92
- Add or Edit IPv6 Pool Dialog Box, page 6-93
- Add or Edit MAC Address Pool Dialog Box, page 6-94
- Add or Edit NET Pool Object Dialog Box, page 6-95
- Add or Edit DHCPv6 Pool Dialog Box, page 6-96

Add or Edit IPv4 Pool Dialog Box

Use the Add or Edit IPv4 Pool dialog box to view, create, or edit IPv4 pool objects.

Navigation Path

Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > IPv4 Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

Related Topics

- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2

Field Reference

Table 6-33 Add IPv4 Pool Object Dialog Box

Element	Description
Name	The object name (up to 64 characters). Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Type	Select whether the pool object is a single IP address or a range of IP addresses.
Address	The IPv4 address of the single host to include in the object.
Start Address End Address	The first and last IP address that define a range of addresses. The start and end addresses must be different, with the start being lower than the end.
Mask	The subnet mask for the IP address or address range.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. **Note** IPv4 Pool objects are always overridable. Clearing this option results in an error. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit IPv6 Pool Dialog Box

Use the Add or Edit IPv6 Pool dialog box to view, create, or edit IPv6 pool objects.

Navigation Path

Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > IPv6 Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

Related Topics

- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2

Field Reference

Table 6-34 Add IPv6 Pool Object Dialog Box

Element	Description
Name	The object name (up to 64 characters). Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Address	The IPv6 address in address/prefix length format to include in the object.
Count	The number of addresses to be included in the pool. Must be between 1 and 16384.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit MAC Address Pool Dialog Box

Use the Add or Edit MAC Address Pool dialog box to view, create, or edit MAC Address pool objects.

Navigation Path

Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > MAC Address Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

Related Topics

- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2

Field Reference

Table 6-35 Add MAC Address Pool Object Dialog Box

Element	Description
Name	The object name (up to 64 characters). Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 6-35 Add MAC Address Pool Object Dialog Box (continued)

Element	Description
Start MAC Address End MAC Address	The first and last MAC address that define a range of addresses. The start and end addresses must be different, with the start being lower than the end.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit NET Pool Object Dialog Box

Use the Add or Edit NET Pool Object dialog box to view, create, or edit Network Entity Title Pool objects.

Navigation Path

Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > NET Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

Related Topics

- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2

Field Reference

Table 6-36 Add NET Pool Object Dialog Box

Element	Description
Name	The object name (up to 64 characters). Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Start NET Address End NET Address	The first and last NET address that define a range of addresses. The start and end addresses must be different, with the start being lower than the end.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-36 Add NET Pool Object Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit DHCPv6 Pool Dialog Box

This dialog box is used to add or edit the DHCPv6 Server Pool. For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature, you can configure the ASA to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients.

Navigation Path

- Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > DHCPv6 Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

OR

- You can access the Add DHCPv6 Pool dialog box from DHCPv6 Pool Selector dialog box: click the Add Row or Edit Row buttons beneath the Available DHCPv6 Pool table. The DHCPv6 Pool Selector dialog box can been accessed from the Server Pool radio button in the Interface IPv6 DHCP section of the IPv6 panel of the Add Interface and Edit Interface dialog box.

Related Topics

- IPv6 Address for Interface Dialog Box, page 46-52
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2

Field Reference

Table 6-37 Add DHCPv6 Pool Dialog Box

Element	Description
Name	The DHCPv6 Pool name should not exceed 200 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.

Table 6-37 Add DHCPv6 Pool Dialog Box (continued)

Element	Description
• Configure parameters on one or more tabs, to provide responses to IR messages to clients. • For each of these tabs, specify the following as appropriate: – DNS/SIP/ NIS/ NISP/ SNTP Server: Enter a server name. Make sure that the IPv6 addresses are in the correct format. For more information on IPv6 address format, see http://www.ietf.org/rfc/rfc2373.txt. – DNS/ SIP/NIS/NISP Domain Name: Enter a domain name. Domain names must begin and end with a digit/letter, only letters, digits and hyphen are allowed as internal characters, labels are separated by a dot.Each label must be up to 63 characters and the entire host name has a maximum of 255 characters. For more information on domain names format, see http://www.ietf.org/rfc/rfc1123.txt. **Note** The import command uses one or more parameters that the ASA obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and in the import command.	
Server tab	(Optional) Specify DNS Server Name and Domain Name.
SIP tab	(Optional) Specify SIP Server Name and SIP Domain Name.
NIS tab	(Optional) Specify NIS Server Name and NIS Domain Name.
NISP tab	(Optional) Specify NISP Server Name and NISP Domain Name.
SNTP tab	(Optional) Specify SNTP Server Name.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring SAML Identity Provider

Beginning with version 4.10, Security Manager enables you to configure Security Assertion Markup Language (SAML) 2.0 based Single-Sign on and Single-Logout for ASA VPN. Single Sign-on Server configuration is no longer supported from ASA version 9.5(2). This has been replaced by SAML Identity Provider.

Security Assertion Markup Language is an XML-based, open-standard data format for exchanging authentication and authorization data between parties, in particular, between an identity provider and a service provider. Identity Provider is a service that can assert a user's identity to another resource. An Identity Provider is responsible for authenticating users in an identity management system. Service Provider is a service that the user wants to access (such as a public or a private web application).

Navigation Path

Select **Manage > Policy Objects** and then select **SAML Identity Provider** from the Object Type Selector.

Adding or Editing SAML Identity Provider

Use the Add or Edit SAML Identity Provider dialog box to add a new SAML Identity Provider or edit an existing row.

Navigation Path

Select **Manage > Policy Objects** and then select **SAML Identity Provider** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Field Reference

Table 6-38 Add or Edit SAML Identity Provider

Element	Description
Name	Enter a name for the SAML Identity Provider, between 4 and 256 characters.
Description	(Optional) Enter a description for the SAML Identity Provider.
Sign In URL	This URL is used for signing into the Identity Provider. It must begin with http:// or https:// (not case sensitive) and the length of the Sign In URL must be less than or equal to 500 characters. The Sign In URL field allows only the following special characters: :, /, *, [,],.
Sign Out URL	(Optional) This URL is used for redirecting to when signing out of the Identity Provider. It must begin with http:// or https:// (not case sensitive) and the length of the Sign Out URL must be less than or equal to 500 characters. The Sign Out URL field allows only the following special characters: :, /, *, [,],.

Element	Description
Base URL	(Optional) This is the clientless VPN's base URL. This URL is used in SAML metadata that is provided to third-party identity providers so that they can redirect end users back to the ASA device. If the Base URL is not configured it is retrieved from the ASA device's hostname and domain name. For example, if the host name is ssl-vpn and domain name is xyz, the Base URL used is https://ssl-vpn.xyz.com. The Base URL must begin with http:// or https:// (not case sensitive) and the length of the Base URL must be less than or equal to 500 characters. The Base URL field allows only the following special characters: :, /, *, [,],. **Note** Either Base URL or Domain Name must be configured in the ASA device to configure SAML.
Identity Provider	Select the Identity Provider from the CA Servers Selector Dialog box. Identity Provider is the service that can assert a user's identity to another resource. An Identity Provider is responsible for authenticating users in an identity management system.
Service Provider	(Optional) Select the Service Provider from the CA Servers Selector Dialog box. Service Provider is the service that the user wants to access (such as a public or private web application).
Request Timeout	(Optional) Enter a value between 1 and 7200. By default there is no SAML timeout.
Enable Signature	(Optional) Enable or disable signature in SAML request. If this is enabled, you must configure the Service Provider. Select the cipher suite for the signature in the Authentication Request drop-down. When you enable the signature, the SHA-256 cipher suite is selected by default. You can change the cipher suite in the Authentication Request drop-down. By default, the Signature is disabled and the Authentication Request drop down is hidden. **Note** You can specify an Authentication Request for a SAML signature, only for ASA 9.8.1 and above.
Enable Internal	(Optional) Enable or disable the Internal flag for SAML Identity Provider. When enabled, the Internal flag identifies the Identity Provider in a private network and the SAML Identity Provider can only be accessed through a WebVPN connection. This also implies that the ASA works as a gateway. By default, the Internal flag is disabled and the Identity Provider can be directly accessed.

Element	Description
Enable Force Re-Authentication	(Optional) Enable or disable Force Re-authentication for SAML Identity Provider. When enabled, the identity provider must authenticate the presenter directly rather than rely on a previous security context. By default, the Force Re-authentication flag is enabled.
Category	(Optional) Select a category between CAT-A to CAT-J.
Allow Value Override per Device	(Optional) If the Allow Value Override per device is selected, edit the Overrides.

Understanding and Specifying Services and Service and Port List Objects

Many policies in Security Manager require that you identify a service to which the policy applies. A service is a protocol and port definition that identifies a particular type of traffic. In many cases, you can specify the service directly in the policy. You can also select service policy objects that define the required services, or use a combination of service objects and policy-specific service designations.

Service objects are convenient because you can create objects to represent the composition of a particular application, or you can model them after the logical organizations that exist on your network, such as a development team or corporate department. There are two types of service policy object:

- Service group—Can contain one or more service, including other service objects. This is the type of service object that was available in all Security Manager 3.x releases.
- Service object—Can contain a single service.

When configuring a policy that requires that you identify a service, you can select or create service objects by clicking the **Select** button next to the Services field. To create a new service from the selection dialog box, click the **Add** button beneath the service list and select a type: group or object. You can also create services from the Policy Object Manager by selecting **Services > Services** from the table of contents and clicking the **Add Object** button and selecting group or object. For information on the specific fields available when creating a service object, see Configuring Service Objects, page 6-103.

Security Manager includes a comprehensive collection of predefined service group objects, including ICMP messages and objects for commonly used services such as HTTP, Syslog, POP3, Telnet, and SNMP. Before using a predefined service group object, you should review the object definition to verify that it conforms to your network implementation. If the predefined object does not meet your needs (for example, if you require different destination ports), you can create a new service object from scratch or based on a copy of an existing object. For more information, see Cloning (Duplicating) Objects, page 6-14.

Whether you are creating a service object or specifying services directly in a policy, you can specify services using the following formats. As you type, Security Manager might prompt you with text-completion options related to your entry. You can select a value from the list and press Enter or Tab. You can enter more than one service by separating services with commas.

- *protocol*, where the protocol is 1-255 or a well known protocol name such as tcp, udp, gre, icmp, and so forth. If you enter a number, Security Manager might convert it to the associated name.
- **icmp/***message_type*/*message_code*, where the message type is 1 to 255 or a well-known ICMP message type name such as echo, and the message code is 0 to 255 (for example, **icmp/unreachable/1** or **icmp/echo-reply**).

- **icmp6/***message_type*/*message_code*, where the message type is 1 to 255 or a well-known ICMP message type name such as echo, and the message code is 0 to 255 (for example, **icmp6/unreachable/1** or **icmp6/echo-reply**).
- **{tcp | udp | tcp&udp}/**{*destination_port_number* | *port_list_object*} where the destination port number is 1-65535 or the name of a port list object. You can enter a range of ports using a hyphen, for example, 10-20. The source port number is the Default Range port list object. The Default Range object includes either all ports (1-65535) or all secure ports (1024-65535), depending on the setting you select in the Policy Objects Page, page 11-66 (select **Tools > Security Manager Administration > Policy Objects**).

  For example, defining a service as tcp/10 means that 10 is the destination port and no source port is defined.

  When you specify ports, you can also use the following special keywords: **lt** (less than), **gt** (greater than), **eq** (equal to), and **neq** (not equal to), followed by a number. For example, **lt 440** specifies all ports less than 440.

Tip To create port list objects, select **Services > Port Lists** in the Policy Object Manager and click the **Add Object** button. For more information, see Configuring Port List Objects, page 6-102.

- **{tcp | udp | tcp&udp}/**{*source_port_number* | *port_list_object*}**/** {*destination_port_number* | *port_list_object*}, where the source and destination port numbers are 1-65535 or the name of a port list object. You can enter a range of ports using a hyphen, for example, 10-20.

  For example, defining a service as tcp/10/20 means that 10 is the source port and 20 is the destination port. If you do not want to specify a destination port, use the Default Range port list object, for example, tcp/10/Default Range.
- (Service groups only) *service_object_name*, which is the name of another existing service object. Specifying other objects lets you nest object definitions. Click **Select** to select a service object or to create a new object.

Note The following ICMP message types which are applicable only on IOS devices are automatically replaced with ASA/PIX/FWSM device supported ICMP message types.

- ICMP-Mobile-Redirect
- ICMP-Host-Unreachable
- ICMP-Network-Redirect
- ICMP-Port-Unreachable
- ICMP-Protocol-Unreachable
- ICMP-Reassembly-Timeout
- ICMP-Redirect
- ICMP-protocol-redirect

Related Topics

- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12

- How Service Objects are Provisioned as Object Groups, page 6-108
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Configuring Port List Objects

Use the Port List dialog box to create, edit, or copy a port list object. Each port list object can contain one or more ports or port ranges (for example, 1-1000 and 2000-2500). Additionally, a port list object can include other port list objects.

You typically use port list objects when defining services, but you can also use them in various policies to identify a port rather than typing in the port number. For more information about using port lists in service definitions, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.

Tip The predefined Default Range port list object includes either all ports (1-65535) or all secure ports (1024-65535), depending on the setting you select in the Security Manager Administration window (select **Tools > Security Manager Administration > Policy Objects** and see Policy Objects Page, page 11-66).

Navigation Path

Select **Manage > Policy Objects**, then select **Services > Port Lists** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Configuring Service Objects, page 6-103

Field Reference

Table 6-39 Port List Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 6-39 Port List Dialog Box (continued)

Element	Description
Ports	The ports or ranges included in the port list object, for example, 443, or 1-1000. You can define a single port, a range of ports, multiple port ranges, or any combination of single ports and ranges. Separate multiple entries with commas. Port values range from 1 to 65535. You can use the following operators to identify ranges: • **gt**—Greater than. For example, gt 1000. • **lt**—Less than. For example, lt 1000. • **eq**—Equals. For example, eq 1000. However, **eq 1000** has the same meaning as simply entering **1000**. • **neq**—Does not equal. For example, neq 1000. If you use this operator, you can include only the neq value in the Ports field. However, you can include port ranges in the object. Thus, if you want to create an object that specifies all ports from 1000-1200 except for 1150, create a port list object for the 1000-1200 range, and another object that specifies neq 1150 and that includes the other port list object.
Port Lists	The other port list objects included in the object, if any. Enter the name of the port lists or click **Select** to select them from a list or to create new objects. Separate multiple entries with commas.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring Service Objects

Use the Add and Edit Service dialog boxes to create or edit service objects. You can create a service object to describe a type of traffic carried by the devices in your network. When creating a service object, you must specify the protocol used by the service.

When you create a service object, you must choose the object type:

- Service Group—Can contain one or more services, including other service objects. This is the type of service object available in all Security Manager 3.x releases.
- Service object—Can contain a single service.

Security Manager provides many predefined service group objects. Before creating an object, scan the list in the Policy Object Manager to see if an existing object fits your needs. Note that although you can duplicate a predefined object, you cannot edit it.

Cisco Security Manager supports the service objects whose definitions are available in the show running configuration. For the predefined objects, the definition would not be available in the show running configuration. Therefore, any policies configured with these predefined objects in ASA device would not be discovered in Cisco Security Manager.

In order to align with device behavior, in Cisco Security Manager version 4.17, support to these predefined objects was introduced. The ASA predefined objects will be discovered, provided, the device is updated with the object whenever a new predefined service object is added in the ASA. This feature is supported for ASA supporting images across all versions.

Note The PPTP predefined object is not supported.

However, Cisco Security Manager does not support activity validation with respect to ASA versions. Also, the ICMP and ICMPv6 predefined objects of Cisco Security Manager are converted to device predefined objects. Hence, any usage of Cisco Security Manager's predefined objects causes negation and re-creation of that policy.

Navigation Path

Select **Manage > Policy Objects**, then select **Services > Services** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type) or right-click a row and select **Edit Object**.

Related Topics

- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Policy Object Manager, page 6-4

Field Reference

Table 6-40 Add and Edit Service Dialog Boxes

Element	Description
Name	The object name. If you are using the object for ASA or PIX devices running software version 8.x, limit the length of the name to 64 characters. For other devices the name can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 6-40 Add and Edit Service Dialog Boxes (continued)

Element	Description
Services (for groups) Service (for objects)	The services to include in this policy object. When creating a Service Group, you can enter more than one service by separating services with commas. When creating a Service Object, you can enter one service only. You can specify services using the following formats. As you type, Security Manager may prompt you with text-completion options related to your entry. If you enter a service that translates directly to a predefined service object, the entry is converted to the predefined object name; for example, TCP/80 is converted to HTTP. • *protocol*, where the protocol is 1 to 255 or a well known protocol name such as tcp, udp, gre, icmp, and so forth. If you enter a number, Security Manager might convert it to the associated name. • **icmp/***message_type/message_code*, where the message type is 1 to 255 or a well-known ICMP message type name such as echo, and the message code is 0 to 255 (for example, **icmp/unreachable/1** or **icmp/echo-reply**). • **icmp6/***message_type/message_code*, where the message type is 1 to 255 or a well-known ICMP message type name such as echo, and the message code is 0 to 255 (for example, **icmp6/unreachable/1** or **icmp6/echo-reply**).
	• **{tcp | udp | tcp&udp}/**{*destination_port_number* | *port_list_object*} where the destination port number can be 1 to 65535, or the name of a port list object. You can enter a range of ports using a hyphen, for example, 10-20. In this instance, the source port number is the Default Range port list object, which specifies the range 1-65535. (See Configuring Port List Objects, page 6-102 for information about creating and editing port list objects.) Whenever you specify ports, you can also use the following special keywords: **lt** (less than), **gt** (greater than), **eq** (equal to), and **neq** (not equal to), followed by a number. For example, **lt 440** specifies all ports less than 440. • **{tcp | udp | tcp&udp}/**{*source_port_number* | *port_list_object*}**/**{*destination_port_number* | *port_list_object*}, where the source and destination port numbers can be 1 to 65535, or the name of a port list object. You can enter a range of ports using a hyphen, for example, 10-20. • (Service groups only) *service_object_name*, which is the name of another existing service object. Specifying other objects lets you nest object definitions. Click **Select** to select a service object or to create a new object.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 6-40 Add and Edit Service Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Note The following ICMP message types which are applicable only on IOS devices are automatically replaced with ASA/PIX/FWSM device supported ICMP message types.

- ICMP-Mobile-Redirect
- ICMP-Host-Unreachable
- ICMP-Network-Redirect
- ICMP-Port-Unreachable
- ICMP-Protocol-Unreachable
- ICMP-Reassembly-Timeout
- ICMP-Redirect
- ICMP-protocol-redirect

How Policy Objects are Provisioned as Object Groups

Object groups are a feature of ASA, PIX, FWSM, and IOS 12.4(20)T+ devices that enable you reduce the size of access rules by grouping objects such as IP hosts, networks, protocols, ports, and ICMP message types. Although the functionality of object groups is similar to the functionality of policy objects in Security Manager, there are several important differences in implementation.

As a result, when deploying policies to a device, it is not always possible to create object groups that are an exact copy of the policy objects that you configured in Security Manager. To take one example, policy object names are unique per object type in Security Manager (that is, you can define a network/host object and a service object with the same name), whereas object groups of all types defined on the device share a single naming scheme. Therefore, if you deploy a network/host object whose name matches an existing service object group on the device, a suffix is added to the name of the network/host object to distinguish it from the service object group.

Note For information about the options available when deploying object groups, see Deployment Page, page 11-13.

Similarly, when discovering policies on a device, it is not always possible to create policy objects that are an exact copy of the object groups that are configured on the device. However, Security Manager preserves as much of the original configuration as possible.

Note For IOS devices, any policy objects that are used by access control list objects are subsequently replaced during deployment by the contents of the object. Object groups used with ACL objects are not preserved, although they are discovered as Security Manager policy objects.

The following sections describe the changes that are made when provisioning policy objects to object groups on the device, or when creating the policy objects when discovering policies on these devices:

- How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups, page 6-107
- How Service Objects are Provisioned as Object Groups, page 6-108

How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups

In most cases, network/host, port list, and service objects can be provisioned as object groups without changing the object name. Table 6-41 on page 6-107 describes how object names are changed when the names cannot be converted directly to object groups on supported devices.

Note The predefined network/host object **any** is not provisioned as an object group.

Table 6-41 How Network/Host, Port List, and Service Objects are Named as Object Groups

Condition	New Name	Examples
Object name includes a space.	Space is replaced with an underscore.	An object named **my object** is provisioned as an object group named **my_object**.
Object name is longer than 64 characters (the maximum supported by object groups).	Name is truncated so that any suffixes required by the object group (such as _TCP or _UDP, or unique numbers, such as _1) can be added while remaining within the 64-character limit.	
Device already has an object group (Protocol/ICMP/Service) with the same name.	A numeric suffix is added to the name, starting from 1.	If you have a network/host object named **West** and the device already has a TCP service object group named **West**, the name of the object group is changed to **West_1** when deployed.
You have already created a network/host object group with the same name.	A numeric suffix is added to the name, starting from 1.	If you have a network/host object and a port list or service object that are both named **West**, the network/host object is deployed as **West** and the port list is deployed as **West_1**.

Note For ASA software versions 8.2 and earlier, if you create an object of the type Network Object on Security Manager, upon discovering the ASA device with an object of the same object name but of type Network-Object Group, Security Manager does not create a new object. Instead, it reuses the existing objects. However, for ASA software versions 8.3 and later, if you create an object of type Network Object on Security Manager, upon discovering the ASA device with an object of the same object name but of type Network-Object Group, Security Manager creates new objects, say, name_1, name_2, and so on. This implies that for Security Manager managing ASA devices that are running the software version 8.2 or earlier, it creates new objects upon upgrading the ASA to version 8.3 or later.

Related Topics

- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- How Service Objects are Provisioned as Object Groups, page 6-108
- How Policy Objects are Provisioned as Object Groups, page 6-106

How Service Objects are Provisioned as Object Groups

The following table describes how Security Manager creates object groups when deploying service objects to supported devices.

Tip For ASA 8.3+ devices, service objects are provisioned using the **object service** command instead of the **object-group** command.

Table 6-42 How Service Objects are Provisioned as Object Groups

Condition	Generated Object Group	Examples

Service object contains the ICMP protocol and ICMP message types.	Generates an ICMP-type object group with the same name as the service object.	Service object service1: **icmp/icmp-echo, 23**
		Object group:
		`object-group icmp-type service1`
		`  icmp-object icmp-echo`
		`  icmp-object 23`

Service object contains only protocols.	Generates a protocol object group with the same name as the service object.	Service object service1: **tcp, gre, 34**
		Object group:
		`object-group protocol service1`
		`  protocol-object tcp`
		`  protocol-object gre`
		`  protocol-object 34`

Service object uses port list objects for both source and destination ports.	Generates service object groups that match the port list objects.	

Table 6-42 How Service Objects are Provisioned as Object Groups (continued)

Condition	Generated Object Group	Examples

Service object contains multiple ports or port ranges, but does not use a port list object for the source ports.	Generates service object group with the name <ObjectName>.src for the source ports.	Service object serv1: **tcp/400,600/23-80**
		Object group:
		`object-group service serv1.src tcp`
		`  port-object eq 400`
		`  port-object eq 600`

Service object contains multiple ports or port ranges, but does not use a port list object for the destination ports.	Generates service object group for the destination ports with the same name as the service object.	Service object serv1: **tcp/400,600/23-80, 566**
		Object group:
		`object-group service serv1 tcp`
		`  port-object range 23 80`
		`  port-object eq 566`
		`object-group service serv1.src tcp`
		`  port-object eq 400`
		`  port-object eq 600`

Service object contains the TCP&UDP protocol and includes defined ports.		Service object serv1: **tcp&udp/400,600/23-80, 566**
		Object group:
		`object-group service serv1 tcp`
		`  port-object range 23 80`
		`  port-object eq 566`
		`object-group service serv1.src tcp`
		`  port-object eq 400`
		`  port-object eq 600`
		`object-group protocol tcp-udp`
		`  protocol-object tcp`
		`  protocol-object udp`

Related Topics

- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- How Network/Host, Port List, and Service Objects are Named When Provisioned As Object Groups, page 6-107
- How Policy Objects are Provisioned as Object Groups, page 6-106

Managing FlexConfigs

FlexConfig policies allow you to configure device commands that are not otherwise supported by Security Manager. By using Flexconfigs, you can extend Security Manager's control over a device configuration and take advantage of new device features before upgrading the product.

FlexConfig policies are made up of FlexConfig objects. These objects are essentially subroutines that can include scripting language commands, device commands, and variables. You can configure an object to be processed prior to applying the Security Manager configuration to a device, or you can have it processed after the configuration. Security Manager processes your objects in the order you specify so that you can create objects whose processing depends on the processing of another object. A FlexConfig policy object's contents can range from a single simple command string to elaborate CLI command structures that incorporate scripting and variables.

Note You can configure Cisco Security Manager either to deploy FlexConfigs only once after creation or modification of a FlexConfig, or to deploy FlexConfigs with each deployment. By default, Cisco Security Manager deploys FlexConfigs one time. If you have FlexConfigs that need to be deployed with each deployment, disable the "Deploy only new or modified Flexconfigs" setting on the **Tools > Security Manager Administration > Deployment page**. After changing this setting, you have to manage one-time FlexConfigs by deleting them after they have been deployed. For more information, see Deployment Page, page 11-13.

Understanding policies and objects is central to understanding and using FlexConfig policy objects. For more information on how Security Manager defines and uses polices, see Chapter 5, "Managing Policies" and for information on how Security Manager defines and uses objects, see Chapter 6, "Managing Policy Objects".

Note With "Deploy only new or modified Flexconfigs" setting on **Tools > Security Manager Administration > Deployment page** enabled, if you have an activity open, with changes, and when you attempt to deploy FlexConfigs, Cisco Security Manager considers the FlexConfig changes specific to that activity alone and not those of other activities. On the other hand, if all activities are submitted and no activity is open, then Security Manager considers the FlexConfig changes specific to the lastly submitted activity that was submitted with changes. Therefore, if you need FlexConfig changes for an activity to reflect during deployment, then ensure the changes are done in a single activity, submitted, and deployed.

The following topics describe FlexConfig policies and policy objects and how to use them:

- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Configuring FlexConfig Policies and Policy Objects, page 7-25

- FlexConfig Policy Page, page 7-36

Understanding FlexConfig Policies and Policy Objects

FlexConfig policy objects are used in FlexConfig policies. They allow you to configure device features that are not otherwise supported by Security Manager, or to otherwise fine-tune your device configurations. These policy objects include device configuration commands, variables, and optionally, scripting language instructions to control processing. FlexConfig objects are essentially programming routines to add content to the device configurations that Security Manager generates.

You can create FlexConfig policy objects from scratch or you can duplicate one of the objects that are included with Security Manager.

FlexConfig policies are simply an ordered list of FlexConfig policy objects. Your objects are processed in the order that you specify.

The following topics help you understand FlexConfig policy objects and by extension, FlexConfig policies. For more information about policy objects in general, see Chapter 6, "Managing Policy Objects".

- Using CLI Commands in FlexConfig Policy Objects, page 7-2
- Using Scripting Language Instructions, page 7-3
- Understanding FlexConfig Object Variables, page 7-5
- Predefined FlexConfig Policy Objects, page 7-19

Using CLI Commands in FlexConfig Policy Objects

The configuration commands that you enter into the FlexConfig Editor are actual CLI commands used to configure devices, such as PIX Firewalls and Cisco IOS Routers. You can include CLI commands that are not supported in Security Manager. You are responsible for knowing and implementing the command according to the proper syntax for the device type. See the command reference for the particular operating system for more information.

When you create a Flexconfig policy object, you determine whether the commands and instructions should be added to the beginning or end of the configuration that is generated from regular Security Manager policies:

- Prepended objects—FlexConfig objects that are processed at the beginning of the configurations. If Security Manager policies configure any of the same commands included in the object, the prepended commands are replaced when configuration files are deployed.
- Appended objects—FlexConfig objects that are processed at the end of the configurations, after all other commands in the configuration file and before the **write mem** command.

  If the appended commands are already configured on the device, the device generates an error when you try to add them again. To resolve this, two workarounds are available:

  – Enter the command that removes the configuration in question as an appended command. For example, if the command is **xyz**, enter the following two lines:

    ```
    no xyz
    xyz
    ```

  – Change the setting that controls the action that the device will take to "warn." This is set under **Tools > Security Administration > Deployment**.

The setting change will affect the behavior of devices for all commands being deployed, not just those designated as appended commands.

Note If you are deploying to a device, you should remove most appended commands after the initial deployment. This is especially true for object groups, where any unbound object group is replaced in the Ending Command section during command generation, then re-sent each time the configuration is deployed to a device. The device displays an error because the firewall device shows that the object group already exists. If you are deploying to a file or AUS, the appended commands should remain.

Using Scripting Language Instructions

You can use scripting language instructions in a FlexConfig policy object to control how the commands in the object are processed. Scripting language instructions are a subset of commands supported in the Velocity Template Engine, a Java-based scripting language that supports looping, if/else statements, and variables.

Security Manager supports all Velocity Template Engine commands except the **include** and **parse** commands. For information about additional supported commands supported, see the Velocity Template Engine documentation.

The following topics provide examples of the most commonly used functions:

- Scripting Language Example 1: Looping, page 7-3
- Scripting Language Example 2: Looping with Two-Dimensional Arrays, page 7-4
- Example 3: Looping with If/Else Statements, page 7-4

Scripting Language Example 1: Looping

A plain old telephone service (POTS) dial peer enables incoming calls to be received by a telephony device by associating a telephone number to a voice port. The following example enables caller ID for a set of POTS dial peers.

Object Body

```
#foreach ($peer_id in ["2", "3", "4"])
    dial-peer voice $peer_id pots
    caller-id
#end
```

CLI Output

```
dial-peer voice 2 pots
caller-id

dial-peer voice 3 pots
caller-id

dial-peer voice 4 pots
caller-id
```

Scripting Language Example 2: Looping with Two-Dimensional Arrays

In this example, a set of phone numbers is associated to voice ports so that incoming calls can be received at a router.

Object Body

```
#foreach ($phone in [ [ "2000", "15105552000", "1/0/0" ], [ "2100",
"15105552100", "1/0/1" ], [ "2200", "15105552200", "1/0/2" ] ] )
    dial-peer voice $phone.get(0) pots
 destination-pattern $phone.get(1)
 port $phone.get(2)
#end
```

CLI Output

```
dial-peer voice 2000 pots
destination-pattern 15105552000
port 1/0/0

dial-peer voice 2100 pots
destination-pattern 15105552100
port 1/0/1

dial-peer voice 2200 pots
destination-pattern 15105552200
port 1/0/2
```

Example 3: Looping with If/Else Statements

In this example, a set of phone numbers is associated to voice ports so that incoming calls can be received at a router. In addition, another set of phone numbers is associated to IP addresses to enable Voice Over IP outgoing calls from the router.

Object Body

```
#foreach ( $phone in [ [ "2000", "15105552000", "1/0/0", "" ],
[ "2100", "15105552100", "1/0/1", "" ],
[ "2200", "15105552200", "", "ipv4:150.50.55.55"]
[ "2300", "15105552300", "", "ipv4:150.50.55.55"] ] )
    dial-peer voice $phone.get(0) pots
     destination-pattern $phone.get(1)
 #if ( $phone.get(2) == "" )
     session target $phone.get(3)
 #else
     port $phone.get(2)
 #end

#end
```

CLI Output

```
dial-peer voice 2000 pots
    destination-pattern 15105552000
 port 1/0/0

dial-peer voice 2100 pots
    destination-pattern 15105552100
 port 1/0/1
```

```
dial-peer voice 2200 pots
    destination-pattern 15105552000
 session target ipv4:150.50.55.55

dial-peer voice 2300 pots
    destination-pattern 15105552300
 session target ipv4:150.50.55.55
```

Understanding FlexConfig Object Variables

Variables in FlexConfig policy objects start with the $ character. For example, in the following line, $inside is a variable:

```
interface $inside
```

There are three types of variables you can use in a FlexConfig policy object:

- Policy object variables—Static variables that reference a specific property. For example, Text objects are a type of policy object variable. They are a name and value pair, and the value can be a single string, a list of strings, or a table of strings. Their flexibility allows you to enter any type of textual data to be referenced and acted upon by any policy object.

  There are three ways to add policy object variables to a FlexConfig policy object. First, move the cursor to the desired location, and then:

  - Right-click and select **Create Text Object**. This command opens a dialog box where you can create a simple single-value text object and assign it a value. When you click OK, the variable is added to the object, and it is added to the list of defined Text objects in the Policy Object Manager window so that you can use it in other objects or edit its definition. For an example of creating simple text variables, see Example of FlexConfig Policy Object Variables, page 7-6.
  - Right-click and select a policy object type from the **Insert Policy Object** sub-menu. These commands open a selector dialog box where you can select the specific policy object that contains the variable that you want to insert. After selecting the policy object, you are presented with the Property Selector dialog box, where you choose the specific property of the object that you want to use and optionally change the name of the variable associated with the property.

    By using this technique, you can add a property from an existing policy object when you know that the property has the value that you want to use. For example, if you want to insert a variable that specifies the RADIUS protocol from the AAA Server Group policy object named RADIUS, you would right-click, select **Insert Policy Object > AAA Server Group**, select RADIUS in the AAA Server Group Selector dialog box, click OK, and then select Protocol in the Object Property field on the AAA Server Group Property Selector dialog box and click OK. The $protocol variable is inserted at the cursor, and the value for the property as defined in the selected object is added to the variables list.
  - Type in a variable name. If you type in a variable, you cannot assign it a value until you click OK on the Add or Edit FlexConfig dialog box. You will be prompted that a variable is undefined, and given the opportunity to define its value. In the FlexConfig Undefined Variable dialog box, you can select the object type of the policy object that contains the desired value, which will prompt you to select the specific policy object and variable. This is essentially identical to the process for inserting policy object variables described above. The technique you use is a matter of personal preference; the end result is the same.

- System variables—Dynamic variables that reference a value during deployment when the configuration is generated. The values are obtained from either the target device or policies configured for the target device. You can declare system variables to be optional in FlexConfig policy objects, which means that the variables do not need to be assigned a value for it to be deployed to the device.

  To insert a system variable into a FlexConfig policy object, move the cursor to the desired location, right-click, and select the variable from the **Insert System Variable** sub-menus. For a description of the available system variables, see FlexConfig System Variables, page 7-7.

- Local Variables—Variables that are local in the looping and assignment derivatives (the **for each** and **set** statements). Local variables get their values directly from the Velocity Template Engine. There is no need to supply values for the local variables.

  To insert a local variable, simply type it in. When you click OK on the Add or Edit FlexConfig dialog box, you will be asked if you want to define the undefined variable. You can click No, or if you click Yes to define other variables, you can leave the object type of the local variable as Undefined.

Example of FlexConfig Policy Object Variables

Using CLI commands and variables, you can create a FlexConfig policy object to name the inside interface and crypto map on a Cisco router:

```
interface $inside
crypto map $mapname
```

The following example shows how to create a FlexConfig policy object that adds these commands and configures the value of $inside as **serial0** and $mapname as **my_crypto**.

When you add the FlexConfig policy object to a device, and the configuration is generated, the following output is created:

```
interface serial0
crypto map my_crypto
```

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **FlexConfigs** from the table of contents. The table in the right pane lists the existing FlexConfig objects.

Step 3 Right-click in the table and select **New Object**. The Add FlexConfig dialog box appears (see Add or Edit FlexConfig Dialog Box, page 7-30).

Step 4 Enter a name and optionally a description for the object.

Tip You can also enter a group name. Groups help you find FlexConfig objects if you create a lot of them. Either type in a group name, or select an existing one from the drop-down list.

Step 5 Keep **Appended** for Type so that the commands are added at the end of the device configuration.

Step 6 Create the content of the object:

a. Click in the FlexConfig edit box (the large white box) and type in **interface** followed by a space.

b. Right-click and select **Create Text Object**.

c. In the Create Text Object dialog box, enter **inside** as the name and **serial0** as the value. Click **OK** to add the variable.

d. Press Enter to move to the next line and type **crypto map** followed by a space.

e. Right-click and select **Create Text Object**.

f. In the Create Text Object dialog box, enter **mapname** as the name and **my_crypto** as the value. Click **OK** to add the variable.

Step 7 Click the **Validate FlexConfig** icon button above the edit box to check the integrity and deployability of the object. If any errors are identified, fix them.

Step 8 Click **OK** to save the policy object. You can now add the object to a device's local or shared FlexConfig policy.

FlexConfig System Variables

System variables reference values during deployment when commands are generated. Security Manager provides a set of defined system variables for you to use in defining FlexConfig policy objects. The values come from the policies you create for the target devices. The values for these variables are required unless otherwise noted. For information about these variables, see the following tables:

- Device system variables—Table 7-1 on page 7-7. For more information about discovering or configuring devices to obtain values for these variables, see Chapter 3, "Managing the Device Inventory".
- Firewall system variables—Table 7-2 on page 7-9. For more information about firewall policies, see Chapter 46, "Managing Firewall Devices" and Chapter 12, "Introduction to Firewall Services".
- Router platform system variables—Table 7-3 on page 7-13. For more information about router policies, see Chapter 61, "Managing Routers".
- VPN system variables—Table 7-4 on page 7-14. For more information about VPN policies, see Chapter 25, "Managing Site-to-Site VPNs: The Basics".
- Remote access system variables—Table 7-5 on page 7-19. For more information about remote access policies, see Chapter 30, "Managing Remote Access VPNs: The Basics".

Table 7-1 Device System Variables (Applying to All Device Types)

Name	Dimension	Description
SYS_DEVICE_IDENTITY	0	Unique device identity for devices managed by a Configuration Engine or Auto Update Server (AUS) as defined on the Tools > Device Properties > General tab. There must be a device identity for devices managed by these servers.
SYS_DOMAIN_NAME	0	DNS domain name as defined on the Tools > Device Properties > General tab. This is not necessarily the same value that is defined in the Platform > Device Admin > Hostname policy.
SYS_FW_OS_MODE	0	Operating system mode of the FWSM or ASA device as defined on the Tools > Device Properties > General tab. Possible values are ROUTER (routed mode), TRANSPARENT, or NOT_APPLICABLE.

Table 7-1 Device System Variables (Applying to All Device Types) (continued)

Name	Dimension	Description
SYS_FW_OS_MULTI	0	Whether the FWSM or ASA is running in single- or multiple-context mode as defined on the Tools > Device Properties > General tab. Possible values are SINGLE, MULTI, or NOT_APPLICABLE.
SYS_HOSTNAME	0	Device hostname as defined on the Tools > Device Properties > General tab. This is not necessarily the same value that is defined in the Platform > Device Admin > Hostname policy.
SYS_IMAGE_NAME	0	Device image name as defined on the Tools > Device Properties > General tab.
SYS_INTERFACE_IP_LIST	1	IP addresses and masks of the interfaces configured in the Interfaces policy. The IP address and mask are in the x.x.x.x/nn format (for example, 10.20.1.2/24). If there are no interfaces defined on the device, no list is returned. Each element in SYS_INTERFACE_NAME_LIST and SYS_INTERFACE_IP_LIST share the same index for the interface. For example, if element 3 in SYS_INTERFACE_NAME_LIST is for Ethernet1, element 3 in SYS_INTERFACE_IP_LIST is the IP address for Ethernet1. If Ethernet1 has no IP address, element 3 in the SYS_INTERFACE_IP_LIST is empty. This variable is optional.
SYS_INTERFACE_NAME_LIST	1	Names of the interfaces on the device configured in the Interfaces policy. If no interfaces are defined on the device, no list is returned. See the explanation above for SYS_INTERFACE_IP_LIST for additional information. This variable is optional.
SYS_MANAGEMENT_IP	0	Management IP address of the device as defined on the Tools > Device Properties > General tab.
SYS_MDF_TYPE	0	Cisco MDF (MetaData Framework) device type, which indicates the device model. This value is displayed on the Tools > Device Properties > General tab, and is determined when you add the device to Security Manager.

Table 7-1 Device System Variables (Applying to All Device Types) (continued)

Name	Dimension	Description
SYS_OS_RUNNING_VERSION	0	Software version of the operating system running on the device as displayed on the Tools > Device Properties > General tab. For example, 12.1, 12.2S, and so on, on an IOS platform. This value is determined when you discover policies from the device.
SYS_OS_TARGET_VERSION	0	Operating system version to be used when generating the device configuration as defined on the Tools > Device Properties > General tab.
SYS_OS_TYPE	0	Operating system for the device as defined on the Tools > Device Properties > General tab. Possible values are IOS, PIX, ASA, FWSM, IPS. You configure this value when you add the device to Security Manager.
SYS_SYS_OID	0	System object ID (SysObjId) of the device, which is determined when you add the device to Security Manager.

Table 7-2 Firewall System Variables

Name	Dimension	Description
SYS_FPM_INPUT_SP	1	FPM policy map names applied on the interface corresponding to the entry in the SYS_FPM_INTERFACE list in the “in” direction. This data is not configured in Security Manager. It is obtained from a router’s running configuration and is used by the IOS_FPM FlexConfig.
SYS_FPM_INTERFACE	1	Interface names. This data is not configured in Security Manager. It is obtained from a router’s running configuration and is used by the IOS_FPM FlexConfig.
SYS_FPM_OUTPUT_SP	1	FPM policy map names applied on the interface corresponding to the entry in the SYS_FPM_INTERFACE list in the “out” direction. This data is not configured in Security Manager. It is obtained from a router’s running configuration and is used by the IOS_FPM FlexConfig.

Table 7-2 Firewall System Variables (continued)

Name	Dimension	Description
SYS_FW_ACL_IN_NAME	1	Names of ACLs applied to interfaces for traffic filtering in the inbound direction. Each element has a one-to-one correspondence with the SYS_INTERFACE_NAME_LIST variable for Cisco IOS routers, PIX Firewalls, Firewall Service Modules, and ASA devices. Configure firewall access rules to generate values for this variable.
SYS_FW_ACL_OUT_NAME	1	Names of ACLs applied to interfaces for traffic filtering in the outbound direction. Each element of this array has a one-to-one correspondence with SYS_INTERFACE_NAME_LIST variable for Cisco IOS routers, PIX Firewalls, Firewall Service Modules, and ASA devices. Configure Access Rules policies to generate values for this variable.
SYS_FW_BRIDGE_INTERFACE_NAMES	1	Names of bridge interfaces. This variable applies only to IOS transparent firewalls. Configure the Firewall > Transparent Rules policies to generate values for this variable.
SYS_FW_ETHERTYPERULE_ACL_NAMES	1	Names of ethertype access-lists applied to interfaces for traffic filtering coming in or going out. Each element of this array has a one-to-one correspondence with the elements in the SYS_FW_ETHERTYPERULE_INTERFACE_NAMES and SYS_FW_ETHERTYPERULE_DIRECTION_NAMES variables. Configure Firewall > Transparent Rules policies to generate values for this variable.
SYS_FW_ETHERTYPERULE_DIRECTION_NAMES	1	Direction that ethertype access-lists are applied. The value is either “in” or “out.” Each element has a one-to-one correspondence with the elements in the SYS_FW_ETHERTYPERULE_ACL_NAMES and SYS_FW_ETHERTYPERULE_INTERFACE_NAMES variables. Configure Firewall > Transparent Rules policies to generate values for this variable.

Table 7-2 Firewall System Variables (continued)

Name	Dimension	Description
SYS_FW_ETHERTYPERULE_INTERFACE_NAMES	1	Interface names to which ethertype access-lists are applied. Each element has a one-to-one correspondence with the SYS_FW_ETHERTYPERULE_ACL_NAMES and SYS_FW_ETHERTYPERULE_DIRECTION_NAMES variables. Configure Firewall > Transparent Rules policies to generate values for this variable.
SYS_FW_INSPECT_IN_NAME	1	Names of Inspect Rules applied to Cisco IOS router interfaces in the inbound direction. Each element of this array has a one-to-one correspondence with the SYS_INTERFACE_NAME_LIST variable for Cisco IOS routers. Configure Inspection Rules policies to generate values for this variable. This variable is optional.
SYS_FW_INSPECT_OUT_NAME	1	Names of Inspect rules applied to Cisco IOS router interfaces in the outbound direction. Each element of this array has a one-to-one correspondence with the SYS_INTERFACE_NAME_LIST variable for Cisco IOS routers. Configure Inspection Rules policies as values for this variable. This variable is optional.
SYS_FW_INTERFACE_HARDWARE_ ID_ LIST	1	Hardware IDs for the device. Configure Interface policies on the device to generate values for this variable. This variable is optional.
SYS_FW_INTERFACE_NETWORK_LIST	1	Interface networks on the device. Configure Interface policies on the device to generate values for this variable.
SYS_FW_INTERFACE_SECURITY_ LEVEL_LIST	1	Interface security levels on the device. Configure Interface policies on the device to generate values for this variable.
SYS_FW_INTERFACE_STATE_LIST	1	Interface states on the device. Configure Interface policies on the device to generate values for this variable.
SYS_FW_INTERFACE_VLAN_ID_LIST	0	VLAN IDs on the device. Configure Interface policies on the device to generate values for this variable.

Table 7-2 Firewall System Variables (continued)

Name	Dimension	Description
SYS_FW_IPV6_ACL_IN_NAME	1	A list of all IPv6 ACLs in the In direction on the device. Configure IPv6 Access Rules policies in the In direction on the device to generate values for this variable.
SYS_FW_IPV6_ACL_OUT_NAME	1	A list of all IPv6 ACLs in the Out direction on the device. Configure IPv6 Access Rules policies in the Out direction on the device to generate values for this variable.
SYS_FW_MPCRULE_TRAFFICFLOW_ TUNNELGROUPNAME	1	Names of tunnel groups specified in Traffic Flow objects. Traffic Flow objects configure **class-map** commands on PIX/ASA devices, and the names of the tunnel groups listed in Traffic Flow objects populate this variable. This variable is used by the ASA_define_traffic_flow_tunnel_group FlexConfig object to create tunnel groups on PIX/ASA devices. This variable is optional.
SYS_FW_MULTICAST_PIM_ ACCEPT_ REG_ROUTEMAP	0	Route-map name used in the **pim accept-register route-map** command. Enter a name for the route-map (Platform > Multicast > PIM > Request Filter), then configure its features using FlexConfig to generate values for this variable. This variable is optional.
SYS_FW_NAT0_ACL_NAMES	1	Names of ACLs used in the **nat** *interface_name* **0 access-list acl_name** command. This variable is optional.
SYS_FW_OSPF_PROCESS_ID_LIST	1	IDs for OSPF routing processes globally configured on PIX Firewalls, Firewall Service Modules, and ASA devices. Configure Platform > Routing > OSPF policies to generate values for this variable.
SYS_FW_OSPF_REDISTRIBUTION_ ROUTE_MAP_LIST	1	Names for the route maps to apply to the OSPF redistribute commands configured on PIX Firewalls, Firewall Service Modules, and ASA devices. Configure Platform > Routing > OSPF policies to generate values for this variable.

Table 7-2 Firewall System Variables (continued)

Name	Dimension	Description
SYS_FW_POLICY_NAT_ACL_NAMES	1	Names of ACLs used in the **policy nat** commands (nat commands with non-0 pool id). Configure NAT (NAT > Translation Rules > Policy NAT) to generate values for this variable. This variable applies to only PIX 6.3(3) and higher, PIX/ASA 7.x, 8.0(x), 8.1(x), and 8.2(x), and FWSM devices. This variable does not apply to Cisco IOS routers. This variable is optional.
SYS_FW_POLICY_STATIC_ACL_NAMES	1	Names of ACLs used in the **policy static** commands that include access lists. Configure NAT 0 (NAT > Translation Rules > Policy NAT) to generate values for this variable. The variable contains the access-list names used by the **nat-0**, **policy nat**, and **policy static** commands. This variable applies to only PIX 6.3(3) and higher, PIX/ASA 7.x, 8.0(x), 8.1(x), and 8.2(x), and FWSM devices. This variable does not apply to Cisco IOS routers. This variable is optional.

Table 7-3 Router Platform System Variables

Name	Dimension	Description
SYS_ROUTER_BGP_AS_NUMBERS_LIST	1	Autonomous system (AS) number of the border gateway protocol (BGP) and exterior gateway protocol (EGP) on the device. Configure Router Platform > Routing > BGP policies to generate values for this variable. This variable is optional.
SYS_ROUTER_EIGRP_AS_NUMBERS_LIST	1	Autonomous system (AS) numbers of the different enhanced Internet gateway routing protocols (EIGRP) and interior gateway protocols (IGP) on the device. Configure Router Platform > Routing > EIGRP policies to generate values for this variable. This variable is optional.

Table 7-3 Router Platform System Variables (continued)

Name	Dimension	Description
SYS_ROUTER_OSPF_PROCESS_ IDS_LIST	1	Open shortest path first (OSPF) interior gateway protocol (IGP) process numbers on the device. Configure Router Platform > Routing > OSPF Process policies to generate values for this variable. This variable is optional.
SYS_ROUTER_QOS_CLASS_MAP_LIST	1	Names of QoS class maps on the device. Configure Quality of Service policies to generate values for this variable. This variable is optional.
SYS_ROUTER_QOS_POLICY_MAP_LIST	1	Names of the QoS policy-maps on the device. Configure Quality of Service policies to generate values for this variable. This variable is optional.

Table 7-4 VPN System Variables

Name	Dimension	Description
Topology Variables related to the VPN in which a device participates. Configure VPNs to generate values for these variables.		
SYS_VPN_TOPOLOGY	1	Virtual private network (VPN) topology type. Possible values are HUB_AND_SPOKE, POINT_TO_POINT, and FULL_MESH.
SYS_VPN_TOPOLOGY_NAME	1	Name of the VPN topology in which the device participates.
SYS_VPN_TOPOLOGY_ROLE	1	Details about the role of the device in the VPN. Possible values are PEER, HUB, and SPOKE.
Devices Variables related to devices in the VPN in which a device participates. Configure VPNs to generate values for these variables.		
SYS_VPN_HOST_NAME	1	Device hostname.
SYS_VPN_LOCAL_PREFIXES	2	Interface and network IP addresses of protected networks.
SYS_VPN_PRIVATE_INTERFACES	2	Private interface names.
SYS_VPN_PRIVATE_TUNNEL_ENDPT_IP	1	Interface tunnel IP address.
SYS_VPN_PUBLIC_INTERFACES	2	Public interface names.

Table 7-4 VPN System Variables (continued)

Name	Dimension	Description
SYS_VPN_TUNNEL_ENDPT_INTERFACE_IP	1	IP address of the VPN endpoint. In IPSec, the endpoint is the VPN interface; in GRE, it is the tunnel source.
SYS_VPN_TUNNEL_ENDPT_INTERFACE_NAME	1	Name of the VPN endpoint. In IPSec, the endpoint is the VPN interface; in GRE, it is the tunnel source.
SYS_VPN_VPNSM_PUBLIC_IFC	2	Export port names for Catalyst 6000 series switches.
Remote Peers Variables related to remote peers in which a device participates. Configure VPNs to generate values for these variables.		
SYS_VPN_REM_PEER_BAK_LOGICAL_PRIVATE_IP	3	Interface tunnel IP addresses of remote peers of failover hubs. This value is used in DMVPN for next hop resolution protocol (NHRP).
SYS_VPN_REM_PEER_BAK_PREFIX	3	Protected networks (interface and network IP addresses) of remote peers of failover hubs.
SYS_VPN_REM_PEER_BAK_PUBLIC_IP	3	Public interface names of remote peers of failover hubs.
SYS_VPN_REM_PEER_BAK_TUNNEL_SRC	3	IP address of the VPN endpoint of remote peers. In IPSec, the endpoint is the VPN interface; in GRE, it is the tunnel source.
SYS_VPN_REM_PEER_DEVICE_NAME	2	Device hostnames of remote peers.
SYS_VPN_REM_PEER_LOGICAL_ PRIVATE_IP	2	Interface tunnel IP addresses of remote peers. This value is used in DMVPN for next hop resolution protocol (NHRP).
SYS_VPN_REM_PEER_PREFIX	3	Protected networks (interface and network IP addresses) of remote peers.
SYS_VPN_REM_PEER_PRIVATE_IP	2	Private interface names of remote peers.
SYS_VPN_REM_PEER_PUBLIC_IP	2	Public interface names of remote peers.
SYS_VPN_REM_PEER_TUNNEL_SRC	2	Tunnel sources (if included in the interface tunnel of remote peers).
IPSec Proposal Variables related to IPSec Proposal policies. For more information, see Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22 and Configuring High Availability in Your VPN Topology, page 25-52. Configure the IPSec Proposal policy to generate values for these variables.		
SYS_VPN_CRYPTO_MAP_TYPE	1	Crypto map type. Possible values are STATIC and DYNAMIC.

Table 7-4 VPN System Variables (continued)

Name	Dimension	Description
SYS_VPN_DYNAMIC_CRYPTO_NAME	1	Dynamic crypto map name.
SYS_VPN_DYNAMIC_CRYPTO_NUM	1	Dynamic crypto map number.
SYS_VPN_STATIC_CRYPTO_NAME	1	Static crypto map name.
SYS_VPN_STATIC_CRYPTO_NAME_BAK	1	Static crypto map name of failover hubs.
SYS_VPN_STATIC_CRYPTO_NUM	2	Static crypto map number.
SYS_VPN_STATIC_CRYPTO_NUM_BAK	2	Static crypto map number of failover hubs.
Preshared Keys Variables related to Preshared Key and IKE Proposal policies. For more information, see Configuring IKEv1 Preshared Key Policies, page 26-48.		
SYS_VPN_IKE_AUTHENTICATION_MODE	1	Authentication method of the IKE policy. Possible values are pre-share, rsa-sig, rsa-encr, dsa-sig. Configure an IKE Proposal policy to generate values for this variable.
SYS_VPN_IKE_PRIORITY	1	Priority number of the IKE policy Configure an IKE Proposal policy to generate values for this variable.
SYS_VPN_NEGOTIATION_MODE	1	Negotiation method. Possible values are MAIN_ADDRESS, MAIN_HOST, and AGGRESSIVE. Configure a Preshared Key policy to generate values for this variable.
GRE Modes Variables related to GRE Modes policies. For more information, see Understanding the GRE Modes Page, page 27-1.		
SYS_VPN_BAK_TUNNEL_IFC	2	Interface tunnel number of remote peers of failover hubs, for example, tunnel0. Configure VPNs to generate values for this variable.
SYS_VPN_SIGP_PROCESS_NUMBER	1	Process number of the interior gateway protocol (IGP). Configure GRE Modes policies to generate values for this variable.

Table 7-4 VPN System Variables (continued)

Name	Dimension	Description
SYS_VPN_SIGP_ROUTING_PROTOCOL	1	Type of secured interior gateway protocol (IGP) used. Possible values are STATIC, OSPF, EIGRP, RIPV2, BGP, and ODR. Configure GRE Modes policies to generate values for this variable.
SYS_VPN_SPOKE_TO_SPOKE_ CONN	1	Whether DMVPN is configured for spoke-to-spoke connectivity. Possible values are true or false. Configure GRE Modes policies to generate values for this variable.
SYS_VPN_TUNNEL_IFC	2	Interface tunnel number of remote peers, for example, tunnel0. Configure VPNs to generate values for this variable.
VRF Variables related to virtual routing and forwarding (VRF). For more information, see Configuring VRF Aware IPsec Settings, page 25-48. Configure VPN VRF settings to generate values for these variables.		
SYS_VPN_VRF_AREA_ID	1	Area ID numbers if the OSPF process number was chosen.
SYS_VPN_VRF_MPLS_INTERFACE_IP	1	Multiprotocol label switching (MPLS) interface IP addresses.
SYS_VPN_VRF_MPLS_INTERFACE_NAME	1	Multiprotocol label switching (MPLS) interface names.
SYS_VPN_VRF_NAME	1	VRF names.
SYS_VPN_VRF_PROCESS_NUMBER	1	Interior gateway protocol (IGP) process numbers.
SYS_VPN_VRF_RD	1	RD values.
SYS_VPN_VRF_ROUTING_PROTOCOL	1	Interior gateway protocol (IGP) values. IGP is used for routing the IPSec aggregator toward the Provider Edge (PE)/Multiprotocol Label Switching (MPLS) network. Possible values are STATIC, OSPF, EIGRP, RIPV2, and BGP.
SYS_VPN_VRF_SOLUTION	1	Virtual routing and forwarding (VRF) solution. Possible values are 1BOX and 2BOX.
CA Variables related to certificate authority policies. For more information, see Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.		

Table 7-4 VPN System Variables (continued)

Name	Dimension	Description
SYS_VPN_CA_NAME	2	Certificate authority (CA) names. Configure PKI policies to generate values for this variable.
EZVPN Variables related to EZVPN. For more information, see Understanding Easy VPN, page 28-1.		
SYS_VPN_EZVPN_GROUP_NAME	2	User group names. Configure User Group policies to generate values for this variable.
Dial Backup Variables related to dial backup configurations. For more information, see Configuring Dial Backup, page 25-40.		
SYS_VPN_RTR_WATCH	1	The rtr/watch number. Configure dial backup to generate values for this variable.
GETVPN Variables related to Group Encrypted Transport (GET) VPN. For more information, see Understanding Group Encrypted Transport (GET) VPNs, page 29-2.		
SYS_GDOI_GROUP_NAME	1	Name of the Group Domain of Interpretation (GDOI) group. Configure the Group Encryption policy to generate values for this variable (Manage > Site-to-Site VPNs > Group Encryption Policy > Group Settings).
SYS_GM_GET_ENABLED_INTF_NAME	1	VPN-enabled outside interface to the provider edge (PE). Traffic originating or terminating on this interface is evaluated for encryption or decryption, as appropriate. Configure group members to generate values for this variable (Manage > Site-to-Site VPNs > Group Members).
SYS_IPSEC_PROFILE_NAME	1	Name of the profile that defines the parameters to be used for IPsec encryption between two group members. Configure the Group Encryption policy to generate values for this variable (Manage > Site-to-Site VPNs > Group Encryption Policy > Security Associations).

Table 7-4 VPN System Variables (continued)

Name	Dimension	Description
SYS_KS_REG_INTERFACE	0	Interface on the key server assigned to handle group domain of interpretation (GDOI) registrations. If no registration interface is specified, GDOI registrations can occur on any interface. Configure key servers to generate values for this variable (Manage > Site-to-Site VPNs > Key Servers).

Table 7-5 Remote Access System Variables

Name	Dimension	Description
SYS_ASA_RA_TUNNEL_GROUP_NAME	2	Tunnel group name for ASA devices.
SYS_ASA_RA_USER_GROUP_NAME	2	Name of the ASA user group.
SYS_EZVPN_RA_DYNAMIC_CRYPTO_ MAP_NAME	1	Dynamic Crypto map name for EZVPN.
SYS_EZVPN_RA_DYNAMIC_CRYPTO_ MAP_SEQ_NUM	1	Dynamic Crypto map number for EZVPN.
SYS_EZVPN_RA_PUBLIC_INTERFACE_PIX	2	External interface names for EZVPN for PIX firewall and ASA devices only.
SYS_EZVPN_RA_STATIC_CRYPTO_ MAP_NAME	1	Static crypto map names for EZVPN.
SYS_EZVPN_RA_STATIC_CRYPTO_ MAP_SEQ_NUM	1	Static crypto map numbers for EZVPN.
SYS_IOS_RA_CA_NAME	1	Certificate authority (CA) names for Cisco IOS devices.
SYS_IOS_RA_PUBLIC_INTERFACE	1	External interface names for Cisco IOS devices.
SYS_IOS_RA_USER_GROUP	1	User group names for Cisco IOS devices.
SYS_IOS_RA_VRF_NAME	1	Virtual routing and forwarding (VRF) names for Cisco IOS devices.

Predefined FlexConfig Policy Objects

Security Manager provides predefined FlexConfig policy objects for you to use. These policy objects have predefined commands and scripting.

Predefined FlexConfig policy objects are read-only objects. To edit these predefined FlexConfig policy objects, duplicate the desired object, make changes to the copy, and save it with a new name. This way, the original predefined FlexConfigs remain unchanged. For lists of these predefined policy objects and further information on each, see the following tables:

- Predefined ASA FlexConfig Policy Objects—Table 7-8 on page 7-22
- Predefined Catalyst FlexConfig Policy Objects—Table 7-7 on page 7-22
- Predefined Cisco IOS FlexConfig Policy Objects—Table 7-8 on page 7-22
- Predefined PIX Firewall FlexConfig Policy Objects—Table 7-9 on page 7-23
- Predefined Router FlexConfig Policy Objects—Table 7-10 on page 7-24

Table 7-6 Predefined ASA FlexConfig Policy Objects

Name	Description
ASA_add_ACEs	Adds an access control entry (ACE) to all access control lists on the device.
ASA_add_EtherType_ACL_ remark	Loops through a list of ethertype access-list names and adds ACEs or remarks to them. The ethertype access list is the same as Transparent Rules for Firewalls in Security Manager. The remarks set by the CLI in this FlexConfig will be shown in the description field of a transparent rule.
ASA_add_IPv6_ACEs	Loops through a list of IPv6 access lists and adds a **deny ip any any log** entry to the end of the ACL.
ASA_command_alias	Creates a command alias named “save” for the **copy running-config** and **copy startup_config** commands.
ASA_copy_image	Copies an image package from a TFTP server to flash.
ASA_csd_image	Provides an ASA Cisco Secure Desktop image. It copies the CSD image from /CSCOpx/tftpboot/*device-hostname* on the Cisco Security Manager server to the device, then configures the CSD image path. Make sure you fill out the device’s hostname in Device Properties. If the image name is different than the default, you can override it in Device Properties > Policy Object Overrides > Text Objects > AsaCsdImageName. Unassign this FlexConfig from the device after the image has been copied and configured.
ASA_define_traffic_flow_tu nnel _group	Defines site-to-site VPN tunnel groups listed in the SYS_FW_MPCRULE_TRAFFICFLOW_TUNNELGROUPNAME system variable. This variable is populated with tunnel group names defined in Traffic Flow objects.
ASA_established	Permits return access for outbound connections through the security appliance. This command works with an original connection that is outbound from a network and protected by the security appliance and a return connection that is inbound between the same two devices on an external host. Uses the **established** command to specify the destination port that is used for connection lookups, which gives you more control over the command and supports protocols where the destination port is known, but the source port is unknown. The **permitto** and **permitfrom** keywords define the return inbound connection.
ASA_FTP_mode_passive	Sets the FTP mode to passive.

Table 7-6 Predefined ASA FlexConfig Policy Objects (continued)

Name	Description
ASA_generate_route_map	Generates a route map to be used by the **pim accept-register route-map** command configured under Platform > Multicast > PIM > Request Filter. Security Manager exports the route-map name used in the **pim** command so that you can configure it as desired.
ASA_IP_audit	Uses the **ip-audit** command to provide the following: • Sets the default actions (alarm, drop, reset) for packets that match an attack signature. • Sets the default actions (alarm, drop, reset) for packets that match an informational signature. • Creates a named audit policy that identifies the actions to take (alarm, drop, reset) when a packet matches a defined attack signature or an informational signature. • Disables a signature for an audit policy. • Assigns an audit policy to an interface.
ASA_MGCP	Identifies a specific map for defining the parameters for Media Gateway Control Protocol (MGCP) inspection.
ASA_no_router_Id	Removes the router ID for each OSPF process.
ASA_no_shut_Intf	Loops through and enables all interfaces on a device.
ASA_privilege	Sets the privilege levels for the **configuration**, **show** and **clear** commands.
ASA_route_map	Defines a route map for each OSPF process redistribution route map name.
ASA_RSA_KeyPair_generation	Resets and generates RSA key pairs for certificates.
ASA_svc_image	Provides an ASA SSL VPN Client image. It copies the SVC image from /CSCOpx/tftpboot/*device-hostname* on the Cisco Security Manager server to the device, then configures the SVC image path. Make sure you fill out the device's hostname in Device Properties. If the image name is different than the default, you can override it in Device Properties > Policy Object Overrides > Text Objects > AsaSvcImageName. Unassign this FlexConfig from the device after the image has been copied and configured.

Table 7-6 Predefined ASA FlexConfig Policy Objects (continued)

Name	Description
ASA_sysopt	Uses the **sysopt** command to provide the following examples: • Ensures that the maximum TCP segment size does not exceed the value you set or that the minimum is not less than a specified size. • Forces each TCP connection to remain in a shortened TIME_WAIT state of at least 15 seconds after the final normal TCP close-down sequence. • Disables DNS inspection that alters the DNS A record address. • Ignores the authentication key in RADIUS accounting responses. • Enables the web browser to supply a username and password from its cache when it reauthenticates with the virtual HTTP server on the security appliance.
ASA_virtual	Configures virtual HTTP and Telnet servers.

Table 7-7 Predefined Catalyst 6500/7600 FlexConfig Policy Objects

Name	Description
Cat6K_ECLB_algorithm	Sets the Ether Channel load balance algorithm for modules.
Cat6K_ECLB_port_mode	Applies an Ether Channel to the Catalyst trunk ports where IPS sensors are plugged in. Make sure the ports are configure in trunk mode.
Cat6K_ECLB_portchannel	Sets the port channel to trunk mode and add trunk-allowed VLANs.
Cat6K_firewall_multiple_vlan_interfaces	Sets multiple VLAN interfaces mode if multiple SVIs need to be provisioned.

Table 7-8 Predefined Cisco IOS FlexConfig Policy Objects

Name	Description
IOS_add_bridge_interface_desc	Loops through a list of bridge interfaces and adds the description “this is a bridge interface.”
IOS_CA_server	Configures a certificate authority server.
IOS_compress_config	Compresses large Cisco IOS configurations.
IOS_config_root_wireless_station	Creates and configures the root radio station for a wireless LAN on Cisco IOS routers such as the 851 and 871.
IOS_console_AAA_bypass	Provides examples of the following scenarios: • Enables the authentication, authorization, and accounting (AAA) access-control model. • Sets AAA at login. • Enables AAA authentication for logins.

Table 7-8 Predefined Cisco IOS FlexConfig Policy Objects (continued)

Name	Description
IOS_Copy_Image	Copies the an SVC image from the Security Manager server to the device, then configures the SVC image path. Unassign this FlexConfig from the device after the image has been copied and configured.
IOS_enable_SSL	Enables SSL.
IOS_FPM	Copies traffic class definition files to a router and applies policy-maps.
IOS_IPS_PUBLIC_KEY	Defines public keys on an IOS IPS device. Public keys are required for Security Manager to perform signature updates.
IOS_IPS_SIGNATURE_CATEGORY	Retires all signatures except those in the ios_ips basic category.
IOS_PKI_with_AAA	Configures a PKI AAA authorization using the entire subject name.
IOS_set_clock	Sets the clock to the current time on the Security Manager server.
IOS_VOIP_advance	Loops through and associates a POTS port number to a telephone number and port or IP address number.
IOS_VOIP_simple	Associates a POTS port number to a telephone number and port number.
IOS_VPN_config_gre_tunnel	Uses VPN variables to configure the GRE tunnel for each VPN in which the device participates.
IOS_VPN_set_interface_desc	Using VPN variables, updates the description of the public interface for each VPN in which the device participates.
IOS_VPN_shutdown_inside_interface	Using VPN variables, shuts down all inside interfaces for each VPN in which the device participates.
IOS_VRF_on_vFW	Configures virtual routing and forwarding (VRF) on virtual firewall interfaces.

Table 7-9 Predefined PIX 6.3 Firewall FlexConfig Policy Objects

Name	Description
PIX6.3_nat0_acl_compiled	Generates a compiled access list for NAT 0 access-control lists.
PIX6.3_policy_nat_acl_compiled	Generates a compiled access list for Policy NAT ACLs
PIX6.3_policy_static_acl_compiled	Generates a compiled access list for Policy Static ACLs.
PIX_VPDN	Configures a virtual private dialup network (VPDN).

Table 7-10 Predefined Router FlexConfig Policy Objects

Name	Description
ROUTER_add_inspect_rules	Loops through and appends inspect rules.
ROUTER_BGP_no_auto_summary	Disables the auto route summary for each BGP process by using the **no auto-summary** sub-command. This FlexConfig policy object uses the list of border gateway protocol (BGP) numbers from the SYS_ROUTER_BGP_AS_NUMBERS_LIST system variable.
ROUTER_BGP_untrusted_info	Uses the **distance bgp 255 255 255** sub-command to make the border gateway protocol (BGP) routing information untrusted for each BGP. This FlexConfig policy object uses the list of BGP numbers from the SYS_ROUTER_BGP_AS_NUMBERS_LIST system variable.
ROUTER_EIGRP_min_cost_routes	Configures traffic to use minimum cost routes when multiple routes have different cost routes to the same destination network. This is done using multi-interface load splitting on different interfaces with equal cost paths. This FlexConfig policy object uses the list of router enhanced interior gateway routing protocol (EIGRP) numbers from the SYS_ROUTER_EIGRP_AS_NUMBERS_LIST system variable.
Router_EIGRP_no_auto_summary	Disables the auto route summary for each router enhanced interior gateway routing protocol (EIGRP) processes by using the **no auto-summary** sub-command. This FlexConfig policy object uses the list of EIGRP numbers from the SYS_ROUTER_EIGRP_AS_NUMBERS_LIST system variable.
ROUTER_interface_prevent_dos _attacks	Prevents denial-of-service (DOS) attacks on all device interfaces. This FlexConfig policy object uses the list of interface names from the SYS_INTERFACE_NAME_LIST system variable.
ROUTER_OSPF_no_router_Id	Removes the router OSPF ID for each OSPF process. This FlexConfig policy uses the list of OSPF IDs from the SYS_ROUTER_OSPF_PROCESS_IDS_LIST system variable.
ROUTER_QoS_Class_Map _description	Sets QoS class map descriptions. This FlexConfig policy object uses the list of router QoS class names from the SYS_ROUTER_QOS_CLASS_MAP_LIST system variable.

Table 7-10 Predefined Router FlexConfig Policy Objects (continued)

Name	Description
ROUTER_QoS_Policy_Map _description	Sets QoS policy descriptions. This FlexConfig policy object uses the list of router QoS policy names from the SYS_ROUTER_QOS_POLICY_MAP_LIST system variable.

Configuring FlexConfig Policies and Policy Objects

You create and manage FlexConfig policy objects in the same way that you create other policy objects. The following topics describe how to create FlexConfig policies and policy objects. For information on other tasks you can perform with FlexConfig policy objects (such as deleting them), see Working with Policy Objects—Basic Procedures, page 6-9.

- A FlexConfig Creation Scenario, page 7-25
- Creating FlexConfig Policy Objects, page 7-28
- Editing FlexConfig Policies, page 7-35

A FlexConfig Creation Scenario

This scenario takes you through the steps to set up Media Gateway Control Protocol (MGCP) for an ASA device using one of the predefined FlexConfig policy objects that are shipped with Security Manager. MGCP is used by the call agent application to control media gateways (devices that convert telephone circuit audio to data packets). Security Manager does not support MGCP configuration, but you can use a FlexConfig policy object to provide a configuration. This illustrates how FlexConfigs enable you to customize, for your network, what is not otherwise supported in Security Manager.

In this scenario, you do the following:

1. Create a policy object by duplicating an existing policy object.
2. Assign the policy object to a device.
3. Preview the configuration to verify that it is correct.
4. Share the policy object with another device.
5. Deploy the configuration to the devices.

You can use this scenario as an example to implement other features by creating copies of and modifying predefined FlexConfig policy objects or by creating your own objects.

Before You Begin

Add two ASA devices to Security Manager for this scenario.

Step 1 Duplicate the FlexConfig policy object by doing the following:

a. Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

b. Select **FlexConfigs** from the table of contents. The table in the right pane lists the existing FlexConfig objects.

c. Right-click ASA_MGCP FlexConfig and select **Clone Object**. The Add FlexConfig dialog box appears (see Add or Edit FlexConfig Dialog Box, page 7-30).

d. Enter a name for the new FlexConfig object, for this example, MyASA_MGCP.

e. Enter a new group name and a description of the object.

Tip The group name and description are optional. We recommend you establish descriptions and groups for objects you create.

f. Click **OK**. The new FlexConfig object appears in the list.

Step 2 Duplicate and edit the $callAgentList text object.

The original ASA_MGCP FlexConfig object uses the variable $callAgentList, which is a text object. The text object is read-only and cannot be edited. Duplicating the text object enables you to edit the duplicate object to apply to your network settings.

a. Select **Text Objects** from the table of contents.

b. Right-click **callAgentList** and select **Clone Object**. The Add Text Object dialog box appears.

c. Edit the name of the text object. For this example change it to mycallAgentList.

d. Double-click the first value in column A and enter the IP address for a call agent in your network. For this example, change the value to 10.10.10.10.

e. Double-click the first value in column B and enter the port number for a call agent in your network. For this example, change the value to 105.

f. Change the IP address and port number values for another call agent. For this example, change the IP address to 20.20.20.20 and the port number to 106. Or, if you have only one call agent in your network, you could remove the second row in the table by decreasing the number in the Number of Rows field. Similarly, if you have *more* than two call agents, you can add rows by increasing the number in this field.

This concept is similar for increasing and decreasing the number of columns by increasing or decreasing the Number of Columns field.

g. Click **OK**. The new text object appears in the list of text objects.

Step 3 Edit the new FlexConfig policy object to use the new variable by doing the following:

a. Select **FlexConfigs** from the table of contents.

b. Double-click MyASA_MGCP. The Edit FlexConfig dialog box appears.

c. Edit $callAgentList to read $mycallAgentList.

d. Click **OK**.

A warning appears that reads: “The following variables are undefined: mycallAgentList Define them now?”

e. Click **Yes** to the warning.

The FlexConfig Undefined Variables dialog box appears with mycallAgentList listed in the Variable Name column.

f. From the Object Type list, select **Text Objects**. The Text Objects window appears.

g. Select **mycallAgentList** from the Available Text Objects list and click **OK**.

h. In the FlexConfig Undefined Variables window, click **OK**.

The mycallAgentList variable appears in the Variables list of the Edit FlexConfig dialog box.

i. In the Edit FlexConfig dialog box, click **OK**.

j. Close the Policy Object Manager window.

Step 4 Assign the new FlexConfig policy object to a device by doing the following:

a. From the Device view, select the device for which you want to set up MGCP.

b. Select **FlexConfigs** from the Policy selector. The FlexConfigs Policy page appears.

c. Click the **Add** button. The FlexConfigs Selector dialog box appears.

d. Select the new MyASA_MGCP FlexConfig policy object and click **>>** to add the policy object to the Selected FlexConfigs column.

You can select multiple policy objects at one time by holding either the Ctrl (for multiple selections) or Shift (for multiple continuous selections) keys while selecting.

e. Click **OK**.

The MyASA_MGCP policy object is added to the Appended FlexConfigs table, because the object is set to be appended to the configuration. You configure FlexConfig policy objects that you want added to the beginning of the configuration as prepended policy objects.

f. Click **Save**.

Step 5 Preview the commands before they are generated and sent to the device by doing the following:

a. From the FlexConfigs Policy page, select the MyASA_MGCP policy object.

b. Click **Preview**.

The commands that are generated with this FlexConfig policy object and the values assigned to the selected device appear. Note the changed values:

```
class-map sj_mgcp_class
  match access-list mgcp_list
  exit

mgcp-map inbound_mgcp
       call-agent 10.10.10.10 105
       call-agent 20.20.20.20 106

       gateway 10.10.10.115 101
       gateway 10.10.10.116 102

  command-queue 150
  exit

policy-map inbound_policy
  class sj_mgcp_class
    inspect mgcp inbound_mgcp
    exit
exit

service-policy inbound_policy interface outside
```

Step 6 If you have additional ASA devices that require MGCP, you can share this policy with them by doing the following:

a. Right-click **FlexConfigs** in the Policy selector and select **Share Policy**.

The Share Policy dialog box appears.

b. Enter a name for the policy and click **OK**. For this example, enter MyShared_ASA_MGCP.

The banner above the FlexConfigs policy now shows that the device is using a shared policy and displays the name of the policy.

c. In the FlexConfigs banner, click the link in the Assigned To field. In this example, the link should be labeled **1 Device**, which indicates that this shared policy is assigned to one device (the device you are viewing).

Clicking the link opens the Shared Policy Assignments dialog box. Using this dialog box, you can select the other devices that should use this policy in the Available Devices list, and click **>>** to add them to the list of devices that are assigned the policy.

d. Click **OK**. The Shared Policy Assignments dialog box closes, and the additional devices you selected are configured to use the shared policy. The link in the banner changes to indicate the number of devices that now use this policy (in this example, **2 Devices**).

Tip You can also share policies from Policy view. Select **View > Policy View**, select FlexConfigs in the policy type selector, select the MyShared_ASA_MGCP policy, click the Assignments tab, select the devices to which you want to assign the policy, click**>>**, and then **Save**.

Step 7 Submit your changes and deploy the configurations to the devices. For information about deploying configurations, see Working with Deployment and the Configuration Archive, page 8-25.

Creating FlexConfig Policy Objects

You can create FlexConfig policy objects to configure features on devices that are not supported by Security Manager. For more information about FlexConfig objects, see Understanding FlexConfig Policies and Policy Objects, page 7-2.

Tip You can also create FlexConfig policy objects when defining policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Before You Begin

Ensure that your commands do not conflict in any way with the VPN or firewall configuration on the devices.

Keep the following in mind:

- Security Manager does not manipulate or validate your commands; it simply deploys them to the devices.
- If there is more than one set of commands for an interface, only the last set of commands is deployed. Therefore, we recommend you not use beginning and ending commands to configure interfaces.
- When editing FlexConfig objects that involve route-maps (for example, OSPF or multicast route-maps), you must define the corresponding access control lists (ACLs) before the route-maps. This is a device requirement. If you do not define ACLs before route-maps, you will get a deployment error.

Related Topics

- A FlexConfig Creation Scenario, page 7-25
- Working with Policy Objects—Basic Procedures, page 6-9
- Creating Policy Objects, page 6-9
- Chapter 5, "Managing Policies"

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager window (see Policy Object Manager, page 6-4).

Step 2 Select **FlexConfigs** from the Policy Object Type selector.

Step 3 Right-click inside the work area and select **New Object**.

The Add FlexConfig Object dialog box appears (see Add or Edit FlexConfig Dialog Box, page 7-30).

Step 4 Enter a name for the object and optionally a description. Other optional informational fields include:

- Group—Select an existing group name or type in a new one. These names can help you identify the use of an object.
- Negate For—If this FlexConfig object is designed to negate another, enter the name of the FlexConfig object whose commands are undone by this object.

Step 5 In the Type field, select whether commands in the object are to be prepended (put at the beginning) or appended (put at the end) of the configurations generated from Security Manager policies.

Step 6 In the object body area, enter the commands and instructions to produce the desired configuration file output. You can type in the following types of data:

- Scripting commands to control processing. For more information, see Using Scripting Language Instructions, page 7-3.
- CLI commands that are supported by the operating system running on the devices to which you will deploy the FlexConfig policy object. For more information, see Using CLI Commands in FlexConfig Policy Objects, page 7-2.
- Variables. You can insert variables using the right-click menu, which allows you to create simple single-value text variables (**Create Text Object**), select variables from existing policy objects (**Insert Policy Object**), or select system variables (**Insert System Variable**). For more information, see Understanding FlexConfig Object Variables, page 7-5.

Tip If you want to remove a variable, select it in the object body and click the Cut button or press the Backspace or Delete key. When you click **OK** to save your changes, the variable is removed from the list of variables.

Step 7 Click the **Validate FlexConfig** icon button above the object body to check the integrity and deployability of the object.

Step 8 Click **OK** to save the object.

Add or Edit FlexConfig Dialog Box

Use the Add or Edit FlexConfig dialog box to create or edit FlexConfig policy objects. FlexConfig objects are small programs that allow you to add configuration commands before or after the configurations generated from Security Manager policies, so that you can extend the abilities of the product to configure your devices. You use these policy objects in FlexConfig device or shared policies.

Before creating FlexConfig policy objects, read the sections in Understanding FlexConfig Policies and Policy Objects, page 7-2.

Navigation Path

Select **Manage > Policy Objects**, then select **FlexConfigs** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Creating FlexConfig Policy Objects, page 7-28
- Editing FlexConfig Policies, page 7-35
- Chapter 7, “Managing FlexConfigs”

Field Reference

Table 7-11 FlexConfigs Editor Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Group	The name of the group of FlexConfig objects to which this object belongs, if any. You can type in a name, or select an existing name from the list. This field is for informational purposes only, and can help you find a FlexConfig object in the FlexConfig Objects page in the Policy Object Manager.
Type	Whether the commands in the object are prepended (put at the beginning) or appended (put at the end) of configurations.
Negate For	The name of the FlexConfig object whose commands are undone in this FlexConfig object. This field is for informational purposes only and does not affect the processing of the object. For example, if FlexConfig A has the command **banner login**, and FlexConfig B has the command **no banner login**, FlexConfig B negates the configuration for FlexConfig A.

Table 7-11 FlexConfigs Editor Dialog Box (continued)

Element	Description
FlexConfig Object Body	
Object Body edit box	The commands and instructions to produce the desired configuration file output. You can type in the following types of data: • Scripting commands to control processing. For more information, see Using Scripting Language Instructions, page 7-3. • CLI commands that are supported by the operating system running on the devices to which you will deploy the FlexConfig policy object. For more information, see Using CLI Commands in FlexConfig Policy Objects, page 7-2. • Variables. You can insert variables using the right-click menu, which allows you to create simple single-value text variables (**Create Text Object**), select variables from existing policy objects (**Insert Policy Object**), or select system variables (**Insert System Variable**). For more information, see Understanding FlexConfig Object Variables, page 7-5.
Undo button	Deletes the previous action.
Redo button	Performs the previously undone action.
Cut button	Deletes the highlighted text and copies it to the clipboard.
Copy button	Copies the highlighted text to the clipboard.
Paste button	Pastes previously cut or copied text.
Find button	Locates the specified text string in the object body.
Validate FlexConfig button	Checks the integrity and deployability of the FlexConfig object.
FlexConfig Object Variables This table lists the variables that are used in the FlexConfig object.	
Name	The name of the variable. Click the cell to edit the name, which also changes the name in the FlexConfig object body.
Default Value	The value to use when one is not provided. Click the cell to edit the value for user-defined variables. You cannot edit system-defined variables. **Note** Except for optional variables, if a default value is not provided, you must provide a value for the variable.
Object Property	The property of the object. The object property name is in the following format: *type.name*.data.*property* where • *Type*—The type of object, for example Text, Network, AAA Server, and so on. • *Name*—The name of the object. • Data—Indicates that the property of the object is data. • *Property*—The property of the data.

Table 7-11 FlexConfigs Editor Dialog Box (continued)

Element	Description
Dimension	The structure of the data in the variable. Possible values are: • 0—scaler (a single string) • 1—one-dimensional array (a list of strings) • 2—two-dimensional table (a table of strings)
Optional	Whether the variable is required to have a value.
Description	A description of the contents of the object. Click the cell to edit the description.

Create Text Object Dialog Box

Use the Create Text Object dialog box as a shortcut to create text objects of dimension 0, which are single-value variables, for use in FlexConfig policy objects. Enter the name of the variable and the value to assign to it. When you click **OK**, the variable is added to the FlexConfig object at the cursor location and it is added to the list of variables for the object.

Navigation Path

In the Add or Edit FlexConfig Dialog Box, right-click in the object body field and select **Create Text Object**.

Tip If you want to create a multiple-value text object, right-click and select **Insert Policy Object > Text Objects**, and click the **Add** button under the available objects list. For more information, see Add or Edit Text Object Dialog Box, page 7-32.

Add or Edit Text Object Dialog Box

Use the Add or Edit Text Object dialog box to create, edit, duplicate, and view text objects. Create a text object if you need textual data to be referenced and acted upon by another policy object that accepts text objects.

Text objects are a type of policy object variable. They are a name and value pair, where the value can be a single string, a list of strings, or a table of strings. You can enter any type of textual data to be referenced and acted upon by FlexConfig policies. For more information about FlexConfigs, see Chapter 7, “Managing FlexConfigs”.

Create the variable by first selecting the dimension: a simple single-value variable (dimension 0), a list of variables (dimension 1) or a table or variables (dimension 2). After you create the desired grid by selecting the dimension and if applicable, the number of rows and columns, enter the data into each cell by first clicking in the cell.

Navigation Path

Select **Manage > Policy Objects**, then select **Text Objects** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Field Reference

Table 7-12 Text Object Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object (up to 1024 characters).
Dimension	The structure of the data in the variable: • 0—scalar (a single string) • 1—one-dimensional array (a list of strings) • 2—two-dimensional table (a table of strings)
Number of Rows	The number of data rows in the variable if the dimension is 1 or 2.
Number of Columns	The number of data columns in the variable if the dimension is 2.
[text field]	The content of the text object. Click the cell and enter the data.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

FlexConfig Undefined Variables Dialog Box

Use the FlexConfig Undefined Variables dialog box to define variables used in the FlexConfig object that have not yet been defined. You can choose from a list of policy object types or add a new policy object to use.

Each row in the table represents a single undefined variable.

Tip You do not need to define local variables, those used by the scripting language for processing control. For more information about variables, see Understanding FlexConfig Object Variables, page 7-5.

Navigation Path

In the Add or Edit FlexConfig Dialog Box, if you enter a variable name but do not define its values, when you click **OK**, Security Manager displays a warning and asks if you want to define the variables. If you click **Yes**, this dialog box is opened.

Field Reference

Table 7-13 FlexConfig Undefined Variables Dialog Box

Element	Description
Variable Name	The name of the undefined variable that you used in the FlexConfig object.
Object Type	The type of policy object that contains the value you want to assign to the variable. For local variables, use the Undefined object type. For variables you want to define, you must select the specific policy object and value within that object to assign to the selected variable. You start by selecting the type of policy object from this list. You are then prompted to select the specific policy object. When you click **OK**, you are prompted to select the specific property within that object that contains the desired value (see Property Selector Dialog Box, page 7-34). When you select the value on the Property Selector dialog box and click **OK**, the value is assigned to the variable.
Object Property	The property of the object. For a detailed explanation, see Add or Edit FlexConfig Dialog Box.
Optional	Whether the variable is required to have a value.

Property Selector Dialog Box

Use the Property Selector dialog box to select the specific property within a selected policy object that you want to assign to a variable within a FlexConfig policy object. The title of the dialog box indicates the type of policy object that you selected (for example, AAA Server Groups Property Selector).

For more information on variables, see Understanding FlexConfig Object Variables, page 7-5.

Navigation Path

- In the Add or Edit FlexConfig Dialog Box, right-click and select a specific policy object group type from the **Insert Policy Object** menu, select a specific policy object when prompted, and click **OK**.
- In the FlexConfig Undefined Variables Dialog Box, select a policy object type from the Object Type field, select a specific policy object when prompted, and click **OK**.

Field Reference

Table 7-14 Property Selector Dialog Box

Element	Description
Object Property	The property of the object that contains the value you want to assign to the variable. For specific information on the properties, see topics related to the configuration of those objects.
Name	The name of variable. This field is not available when you are defining undefined variables.
Description	An optional description of the variable. This field is not available when you are defining undefined variables.

Editing FlexConfig Policies

You can assign FlexConfig policies to devices using either Device view or Policy view (for shared policies) by selecting **FlexConfigs** from the policy selector. You can deploy configurations containing these policies as you would deploy any configuration generated by Security Manager. For a scenario that takes you through setting up a FlexConfig policy object and creating a shared FlexConfig policy, see A FlexConfig Creation Scenario, page 7-25.

Note With "Deploy only new or modified Flexconfigs" setting on **Tools > Security Manager Administration > Deployment page** enabled, if you have an activity open, with changes, and when you attempt to deploy FlexConfigs, Cisco Security Manager considers the FlexConfig changes specific to that activity alone and not those of other activities. On the other hand, if all activities are submitted and no activity is open, then Security Manager considers the FlexConfig changes specific to the lastly submitted activity that was submitted with changes. Therefore, if you need FlexConfig changes for an activity to reflect during deployment, then ensure the changes are done in a single activity, submitted, and deployed.

When you edit a FlexConfig policy, you can perform the following actions:

- **Add FlexConfig objects**—To add a FlexConfig object to a policy, click the Add icon button and select the desired object. You can also create new objects from the object selector dialog box. The objects are added to the prepended or appended list depending on how the objects themselves are defined.
- **Remove FlexConfig objects**—If you no longer want to include an object in the policy, select it and click the Remove icon button. This action removes the object from the policy, but it does not delete the object from Security Manager. For information on deleting objects, see Deleting Objects, page 6-16.
- **Change the order of the objects**—Objects are processed in the order you specify. If an object depends on the processing of another object, it is important that you order them correctly. Select the object whose order you want to change and click the Up or Down arrow buttons until the object is in the desired location.

  When changing the order of FlexConfig objects that involve route-maps (for example, OSPF or multicast route-maps), make sure that the corresponding access control lists (ACLs) are defined before the route-maps. This is a device requirement. If you do not define ACLs before route-maps, you will get a deployment error.
- **Change the values assigned to the variables used in a policy object**—If you want to configure a variable with a different value for a particular device, creating a device-level override for the object, select the object and click **Values**. In the Values Assignment dialog box, click in the Values cell to change the value. For more information, see Values Assignment Dialog Box, page 7-37.
- **Preview the CLI that will be generated for a policy object**—In Device view, you can view the CLI that will be generated for a policy object by selecting the object and clicking **Preview**. This is especially useful for checking that the CLI commands generated are what you intend to implement on the device.

  Note During deployment, when the FlexConfig policy objects are compiled on the Security Manager server, the correct system variable values and settings are used to generate commands. However, because the Preview function does not have access to these values the way it normally would during deployment, it might not display some CLI commands. In addition, because the Preview function generates CLI commands on the client, some macros used in FlexConfig policy objects reflect client settings instead of server settings.

Related Topics

- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Creating FlexConfig Policy Objects, page 7-28
- Chapter 5, "Managing Policies"
- Chapter 8, "Managing Deployment"

FlexConfig Policy Page

Use the FlexConfig Policy page to create FlexConfig policies. FlexConfig policies contain ordered lists of FlexConfig policy objects, which are subroutines that allow you to extend the ability of Security Manager to configure your devices. For more information on FlexConfig policy objects, see Understanding FlexConfig Policies and Policy Objects, page 7-2.

Navigation Path

- (Device view) Select **FlexConfigs** from the Policy selector.
- (Policy view) Select **FlexConfigs** from the Policy Type selector and select an existing policy or click the Create a Policy button to create a new one.

Related Topics

- Creating FlexConfig Policy Objects, page 7-28
- Chapter 7, "Managing FlexConfigs"

Field Reference

Table 7-15 FlexConfigs Policy Page

Element	Description
Prepended FlexConfigs	The FlexConfig policy objects that are added to the beginning of the configuration. The objects are processed in the order shown.
Appended FlexConfigs	The FlexConfig policy objects that are added to the end of the configuration. The objects are processed in the order shown.
Values button	Click this button to view, modify, or validate the values assigned to the variables used in the selected FlexConfig policy object using the Values Assignment Dialog Box, page 7-37.
Preview button (Device view only.)	Click this button to view the CLI commands that will be generated for the selected FlexConfig policy object. In Policy view, you can preview CLI by first clicking **Values**, selecting a device in the Values Assignment dialog box, and clicking **Preview**.
Up/Down arrow buttons	Click these buttons to move the selected object up or down in the list. The objects are processed in the displayed order, so it is important that an object whose processing depends on the processing of another object comes after the object it depends on.
Add button	Click this button to add a FlexConfig policy object to the policy. The object itself defines whether it will be added to the prepended or appended list. You can create new FlexConfig objects or select existing ones.

Table 7-15 FlexConfigs Policy Page (continued)

Element	Description
Edit button	Click this button to edit the selected FlexConfig policy object. Your changes affect all devices that use the edited object; your changes are not local policy object overrides for the device. **Note** If you selected a predefined FlexConfig policy object packaged with Security Manager, or an object for which you do not have edit permission, you are allowed only to view the object definition.
Remove button	Click this button to remove the selected object from the policy The object is not deleted from Security Manager; it is simply removed from the FlexConfig policy.

Values Assignment Dialog Box

Use the Values Assignment dialog box to view the variables used in a FlexConfig policy object, validate the object, or preview the CLI generated from the object. For more information, see Understanding FlexConfig Object Variables, page 7-5.

Navigation Path

Select an object and click **Values** from the FlexConfig Policy Page, page 7-36.

Field Reference

Table 7-16 Values Assignment Dialog Box

Element	Description
Assigned Devices (Policy view only)	The devices to which the shared FlexConfig policy has been assigned. Select the device for which you want to display variable values.
Name	The name of the variable.
Value	The value to use for the variable. To change the value, double-click the cell. When you change this value, Security Manager creates a device-level override for the policy object. If the policy object is configured so that its values cannot be overridden, you cannot edit the value. If there is no default value for the variable, you must provide a value unless it is an optional variable.
Default Value	The value assigned to the variable in the policy object. Double-click this cell to view the definition of the policy object that defines the variable's value.
Override	Whether you can override the value of the variable. You can edit the value of only those variables that have a checkmark in this column.
Object Property	The property of the object. For a detailed explanation, see Add or Edit FlexConfig Dialog Box.

Table 7-16 Values Assignment Dialog Box (continued)

Element	Description
Dimension	The structure of the data in the variable: • 0—scalar (a single string) • 1—one-dimensional array (a list of strings) • 2—two-dimensional table (a table of strings)
Optional	Whether the variable value can be empty.
Description	A description of the variable.
Validate button	Click this button to validate the Velocity Template Language syntax and make sure that all required variables have values, that variables do not start with SYS_, and that referenced policy objects exist.
Preview button	Click this button to display the generated CLI commands for the selected FlexConfig policy object.

FlexConfig Preview Dialog Box

Use the FlexConfig Preview dialog box to view the generated CLI commands based on the variables of the selected object defined in the FlexConfig policy.

Navigation Path

To open the FlexConfig Policy Preview dialog box, do one of the following:

- In the Values Assignment Dialog Box, click **Preview**. In Policy view, you must first select a device.
- (Device view) Select a device and click **FlexConfig** (see FlexConfig Policy Page, page 7-36). Select an object in the FlexConfig policy and click **Preview**.

Troubleshooting FlexConfigs

Problem: When adding a FlexConfig using the Cisco Security Manager client, you might receive the following error message:

```
Syntax Error: Failed to setup Velocity Engine to validate syntax.
```

This problem is due to administrator privilege rights on Microsoft Windows. In Microsoft Windows Vista and Microsoft Windows 7, Security Manager requires administrator privileges in order to use the FlexConfig feature.

Solution: To resolve this issue, launch the Security Manager client with administrator privileges in either of the following ways:

- To launch the Security Manager client with administrator privileges, right-click the **Configuration Manager** shortcut and select **Run as administrator**.
- To permanently enable administrator privileges for the Security Manager client, right-click the **Configuration Manager** shortcut and select **Properties**. On the Compatibility tab, select **Run this program as an administrator**, and then select **OK**.

Problem: When using FlexConfigs to deploy to an ASA firewall the following two commands in one job: **reload in x noconfirm**, **reload cancel**, you will receive the following error message:

```
An error response from the device prevented successful completion of this
operation. The device provided the following description: reload cancel No
reload is scheduled
```

Unfortunately, deployment always fails due to fact that both commands are pushed too fast, such that the **reload cancel** is sent before the reload schedule is activated on the device.

Solution: To work around this problem, the commands must be sent in two separate deployments created manually.

Problem: When a FlexConfig is assigned and deployed to a device, the FlexConfig is sometimes shown in subsequent Full configuration previews even after the FlexConfig is deleted from a device.

Solution: No workaround is needed. The FlexConfig will not be included in the deployment because only the delta configuration is pushed to the device during deployment.

Managing Deployment

The settings and policies you define in Security Manager must be deployed to your devices so that you can implement them in your network. The steps you take to deploy configurations to devices depend on whether you are using Workflow mode or non-Workflow mode.

Although non-Workflow mode is the default mode of operation for Security Manager, you can use Workflow mode if your company requires it. For more information, see Workflow and Activities Overview, page 1-20.

The following topics provide information about deploying configurations to devices, in each workflow mode:

- Understanding Deployment, page 8-1
- Working with Deployment and the Configuration Archive, page 8-25
- Rolling Back Configurations, page 8-63

Understanding Deployment

A deployment job defines how configuration changes are sent to devices. In a deployment job, you can define several parameters, such as the devices to which you want to deploy configurations and the method used to deploy configurations to devices. You can also create deployment schedules to automatically spawn deployment jobs at regular intervals.

The following topics will help you better understand and use deployment jobs:

- Overview of the Deployment Process, page 8-1
- Deployment in Non-Workflow Mode, page 8-3
- Deployment in Workflow Mode, page 8-5
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Handling Device OS Version Mismatches, page 8-13

Overview of the Deployment Process

Broadly speaking, deployment is a three-step process, as described in the following table.

Table 8-1 Overview of the Deployment Process

Steps	Deployment Steps
Step 1	Security Manager obtains the current configuration for the device and compares it to the latest saved policies for the device in Security Manager. What Security Manager considers to be the current configuration depends on the type of device, the deployment method, and the settings for deployment preferences. These are the possible sources and the conditions under which they are used: • Obtain the running configuration from the device. The running configuration is used when deploying to the device *unless* the deployment method is AUS, TMS, or CNS. You can force Security Manager to use Configuration Archive by selecting **When Deploying to Device Get Reference Config from: Config Archive** as the deployment preference (select **Tools > Security Manager Administration**, then select **Deployment**). • Obtain the last full configuration from the Security Manager Configuration Archive. The Configuration Archive is used when: – Deploying to file, unless you select **When Deploying to File Get Reference Config from: Device** as the deployment preference. – The deployment method is TMS or CNS. – The device is not managed by Security Manager. – Deploying to a device if uploading the configuration from the device failed. Configuration Archive is used as a backup to obtaining the configuration from the live device. – You preview configurations. • Use the factory default configuration. The factory default configuration is used with PIX or ASA devices if you use the AUS deployment method. It is used for deployment and for configuration preview.
Step 2	Security Manager builds a delta configuration that contains the commands needed to update the device configuration to make it consistent with the assigned policies. It also builds a full device configuration.
Step 3	If you are deploying to the device, Security Manager deploys either the delta configuration or the full configuration, depending on which deployment method you use. If you are deploying to file, Security Manager creates two files: *device_name*_delta.cfg for the delta configuration, and *device_name*_full.cfg for the full configuration. In both cases, the configurations are also added to Configuration Archive. These are the actions based on deployment method: • SSL (HTTPS), SSH, or Telnet—Security Manager contacts the device directly and sends the delta configuration to it. • Auto Update Server (standalone or running on Configuration Engine) for PIX and ASA devices—Security Manager sends the full configuration to Auto Update Server, where the device retrieves it. The delta configuration is not sent. • Configuration Engine for IOS devices—Security Manager sends the delta configuration to Configuration Engine, where the device retrieves it. • TMS—Security Manager sends the delta configuration to the TMS server, from which it can be downloaded to an eToken to be loaded onto the device.

During deployment, if Security Manager determines that the configuration on the device differs from the last-deployed configuration, Security Manager overwrites the changes by default. You can control this behavior using the deployment preferences; select **Tools > Security Manager Administration**, then select **Deployment**, and look for the **When Out of Band Changes Detected** setting. You can also control this for a specific deployment job by editing the deployment method for the job.

If you make changes to the device configuration outside of Security Manager, you have two choices for bringing those changes into Security Manager:

1. You can rediscover policies on the device, in which case all policies for the device become local policies, and any assignments of shared policies to the device are removed.
2. You can make the required changes in Security Manager and redeploy them to the device. During deployment, do not select the option to force an error if out-of-band changes are found on the device. This is the recommended approach.

For more information on how out-of-band changes affect deployment, see Understanding How Out-of-Band Changes are Handled, page 8-12.

After configurations are deployed, you should make changes only through Security Manager for configurations that Security Manager controls. This varies based on operating system. For IPS devices, Security Manager controls the entire configuration. For IOS, ASA, PIX, and FWSM devices, you have more control over which aspects of the device configuration Security Manager controls. If you do not create policies for a feature in Security Manager, such as routing policies, Security Manager does not control those features on the device. If you do create policies for these features, Security Manager overwrites the settings on the device with the settings you defined in Security Manager. Through administration settings, you can control the types of policies that will be available for these devices, thereby preventing Security Manager from displaying or changing policies for these features. To see the available features and control whether they are available for management in Security Manager, select **Tools > Security Manager Administration**, then select **Policy Management**. Security Manager does manage VPN-related policies.

Related Topics

- Deployment in Non-Workflow Mode, page 8-3
- Deployment in Workflow Mode, page 8-5
- Deployment Page, page 11-13
- Policy Management Page, page 11-64

Deployment in Non-Workflow Mode

These topics help you understand deployment in non-Workflow mode:

- Deployment Task Flow in Non-Workflow Mode, page 8-3
- Job States in Non-Workflow Mode, page 8-4

Deployment Task Flow in Non-Workflow Mode

The deployment task flow in non-Workflow mode consists of three simple steps:

1. **Create the job:** A deployment job is created for you when you do one of the following:
   - Click the **Submit and Deploy Changes** button on the main toolbar, or select **File > Submit and Deploy**.

Note These options are not available when Ticket Management is enabled.

- Select **File > Deploy**.
- Select **Manage > Deployments** and click **Deploy**.

2. **Define the job:** You specify parameters, such as the devices to which you want to deploy the configurations and whether you want to deploy directly to the devices or to a file.

   During this step, you can also preview configurations and compare them to the previously deployed configurations or the configuration currently running on the device.

   Note Devices selected for one job cannot be included in any other job. This measure ensures that the order in which policies are deployed is correct. However, you can include devices that are specified in deployment schedules.

3. **Deploy the job:** Deploying the job sends the generated CLI to devices, either directly or through an intermediary transport server (such as AUS, CNS, or TMS) or to output files. You select the destination (device or file) when defining a job. The transport server is specified in the device properties. For more details about defining deployment methods and transport servers, see Understanding Deployment Methods, page 8-8.

Job States in Non-Workflow Mode

In non-Workflow mode, the Status column on the Deployment Manager window lists the state of each job. The following table lists and describes all possible job states in non-Workflow mode. For more details, see Deployment Manager Window, page 8-16.

Table 8-2 Job States in Non-Workflow Mode

State	Description
Deployed	Configurations for all the devices in the job were successfully deployed to the devices or to configuration files. Devices in the job can now be included in another job.
Deploying	Configurations generated for the job are being deployed to the devices or to a directory on the Security Manager server. You can monitor the job progress in the Deployment Manager window if the Deployment Status window is not already open.
Aborted	The job was manually halted. Devices in the job can now be included in another job.
Failed	The deployment to one or more devices in the job failed. Devices in the job can now be included in another job.
Rolling Back	Security Manager is in the process of reverting to and deploying previous configurations for the devices within the deployment job. You can abort a job that is in the Rolling Back state.
Rolled Back	Security Manager has successfully reverted to and deployed previous configurations for the devices within the deployment job.

Deployment in Workflow Mode

These topics help you understand deployment in Workflow mode:

- Deployment Task Flow in Workflow Mode, page 8-5
- Job States in Workflow Mode, page 8-6
- Deployment Job Approval, page 8-7
- Deployment Jobs and Multiple Users, page 8-8

Deployment Task Flow in Workflow Mode

The following is a typical task flow in Workflow mode (see Figure 8-1):

1. **Create the job:** Before you deploy configurations to your devices, you must create a deployment job.
2. **Define the job:** When you create a job, you specify parameters, such as the devices to which you want to deploy the configurations, whether you want to deploy directly to the devices or to a file, and when you want the job to take place.
3. **Submit the job:** In some organizations, before jobs can be deployed, they must be approved by a separate user with the appropriate permissions. In this case, Workflow mode is enabled *with* a deployment job approver, and you must submit the job to this user for review. The user reviews the job and either approves or rejects it.
4. **Approve or reject the job:** If you are working in Workflow mode with a deployment job approver, the approver reviews it, and can then either approve or reject the job. If the job is approved, the submitter can then deploy the job. If the job is rejected, the submitter can discard the job and start over or modify the job and resubmit it.

   If you are working in workflow mode without an approver, you can approve the job yourself.
5. **Deploy the job:** Deploying the job sends the generated CLI to either devices, intermediary transport servers (such as AUS, CNS, or TMS), or files. You select the destination (device or file) when defining the job. The transport server is specified in the device properties. For more details about defining deployment methods and transport servers, see Understanding Deployment Methods, page 8-8.

For descriptions of job states (shown in red in Figure 8-1), see Job States in Workflow Mode, page 8-6.

Figure 8-1 Deployment Task Flow in Workflow Mode

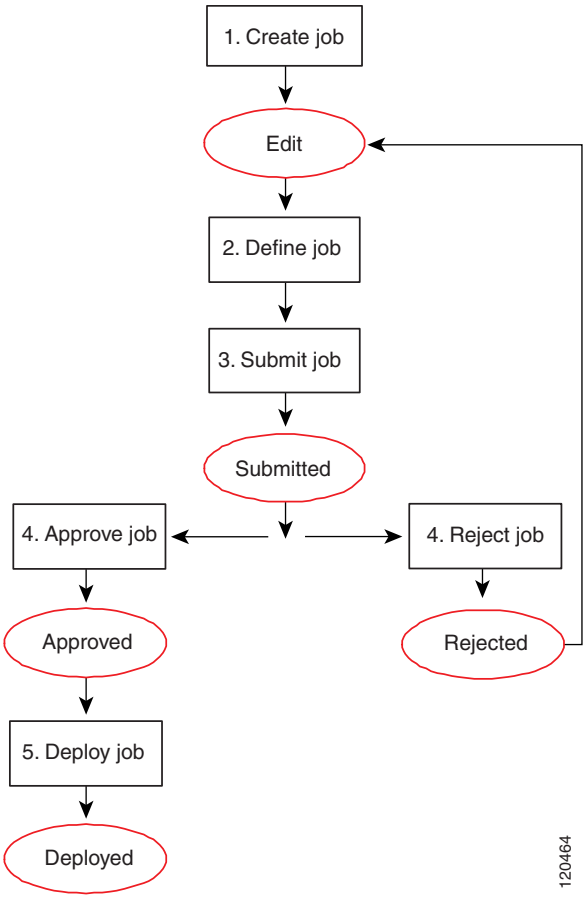

Job States in Workflow Mode

In Workflow mode, the Status column in the Deployment Manager window lists the state of each job. The following table lists and describes all possible job states. For more details about the Deployment Manager window, see Deployment Manager Window, page 8-16.

Table 8-3 Job States in Workflow Mode

State	Description
Edit	The job was created, but it is not currently being edited. The job can be opened, approved (in auto-approval mode), or discarded while it is in the Edit state.
Edit-In Use	The job is open for editing. The job can be closed, approved, discarded, or submitted while it is in the Edit Open state.
Submitted	The job was submitted for review. It can be viewed but not edited while it is in the Submitted state. The job can be opened for viewing, discarded, rejected, or approved while it is in the Submitted state. This state occurs only when Workflow mode is enabled with deployment job approval required.

Table 8-3 Job States in Workflow Mode (continued)

State	Description
Approved	The job was approved and is ready to be deployed. The job can be deployed while it is in the Approved state.
Rejected	The job was rejected. You can open the job for editing or discard the job while it is in the Rejected state. This state occurs only when Workflow mode is enabled with deployment job approval required.
Discarded	The job was discarded. No further changes to the job are not allowed. The job remains in the Deployment table showing a Discarded state until it is purged from the system. Devices in the job can be included in another job.
Deployed	Configurations for all the devices in the job were successfully deployed to the devices or to configuration files. Devices in the job can now be included in another job.
Deploying	Configurations generated for the job are being deployed to the devices or to a directory on the Security Manager server. You can monitor the job progress in the Deployment Manager window.
Aborted	The job was manually halted. Devices in the job can now be included in another job.
Failed	The deployment to one or more devices in the job failed. Devices in the job can now be included in another job.
Scheduled to run at [date]	The job is scheduled to be deployed at the date and time specified.
Rolling Back	Security Manager is in the process of reverting to and deploying previous configurations for the devices within the deployment job. You can abort a job that is in the Rolling Back state.
Rolled Back	Security Manager has successfully reverted to and deployed previous configurations for the devices within the deployment job.

Deployment Job Approval

By default, Security Manager operates in non-Workflow mode; deployment jobs are handled behind the scenes and the user does not need to be aware of jobs or their approval. When using Workflow mode, you can choose to operate with or without a deployment job approver.

If you choose to operate without an approver, you have the permissions to define and approve jobs.

If your organization requires a different person with higher permissions to approve deployment of new or changed configurations to devices, use Workflow mode with a deployment job approver. When using Workflow mode with a deployment job approver, the job must be reviewed by a person with the appropriate permissions to approve or reject the job. This approval process helps to ensure that no inappropriate configurations reach the network devices and that deployment jobs are scheduled effectively.

Note You enable and disable deployment job approval under Tools > Security Manager Administration > Workflow. For more information, see Workflow Page, page 11-75.

Deployment Jobs and Multiple Users

Only one user can define or change parameters or devices within an individual deployment job at one time. However, multiple users can work on the same deployment job in sequence: if a deployment job is closed, another user can open it and make changes to it. Multiple users can work in parallel on different deployment jobs.

Including Devices in Deployment Jobs or Schedules

When you create a deployment job or schedule, you select the devices to include in it. The inclusion of a device influences how the device can be used in other jobs or schedules. When you select a device for a specific job, it cannot be selected for any other job until the original job is deployed, rejected (in Workflow mode), discarded, or aborted. This mechanism prevents two or more people from deploying changes to the same device at the same time and ensures that policies are deployed to devices in the correct order.

However, a device can be part of a deployment schedule and still be selected for specific deployment jobs. While a deployment job is running, the device is locked. The device cannot be included in other jobs while the deployment job is running.

When you create a deployment job, Security Manager displays the devices on which policy changes were made but were not yet deployed. You can deploy to these devices, and you can select additional devices for the job. Although you can add as many devices to a deployment job as you desire (there is no limitation), as a practical matter, you should limit the number of devices per job. The deployment job might fail if you select a large number of devices or several devices that have large configuration files. If you encounter deployment failures, resubmit the job with fewer devices selected.

For VPNs, Security Manager must generate commands for devices that are affected by the policies defined for the devices you select for the job. So, if you select a device that is part of a VPN, Security Manager adds the other relevant devices to the job. For example, if you define a tunnel policy on a spoke, and you select the spoke for the job, Security Manager adds the spoke's assigned hub to the job. During job generation, Security Manager generates commands for both peers so that the VPN configuration is complete and the tunnel can be established. If you deselect one of the devices associated with the VPN, Security Manager warns that removing the device might result in the VPN not functioning property.

Understanding Deployment Methods

Security Manager lets you deploy configurations to devices using three main methods: deploying directly to the device, deploying to a configuration file (which you must then manually apply to the device), and deploying to an intermediate server (which is treated like deploying directly to the device). The system default deployment method is to deploy directly to the device.

When you add devices to Security Manager, you select the deployment method to be used by that device. This determines the method used for deploying to the device (instead of a file). When you create a deployment job, an additional deployment method default applies to the job as a whole, which determines whether deployment creates configuration files or whether it sends the configuration to the device using the method selected for the device. You control this default in the administration settings (select **Tools > Security Manager Administration**, then select **Deployment**; see Deployment Page, page 11-13). When you create a deployment job, you can also change whether the deployment is to a file or to the device for each device by clicking **Edit Deploy Method** in the Create Job window. If you are using non-Workflow mode, see Deploying Configurations in Non-Workflow Mode, page 8-28. If you are using Workflow mode, see Creating and Editing Deployment Jobs, page 8-35.

The method you choose to use depends on the processes and procedures of your organization and the transport protocols supported by a particular type of device. If you are using Configuration Engine (CNS) or Auto Update Server (AUS), use those deployment methods. You must use one of these for devices that use dynamic IP addresses. Otherwise, for devices with static IP addresses, use SSL (HTTPS) for IOS, PIX, ASA, IPS, and standalone FWSM devices, and SSH for FWSM through the Catalyst chassis. If you are using a Token Management Server (TMS) for some devices, you can also use that method with Security Manager.

The following topics describe the deployment methods in more detail:

- Deploying Directly to a Device, page 8-9
- Deploying to a Device through an Intermediate Server, page 8-10
- Deploying to a File, page 8-11
- Understanding How Out-of-Band Changes are Handled, page 8-12

Deploying Directly to a Device

If you choose to deploy directly to a device, Security Manager uses the transport protocol defined in the device properties for the device (right click the device, select **Device Properties**, and click **General**). The protocol is typically the default protocol defined in the Device Communication page in the Security Manager Administration settings (see Device Communication Page, page 11-21). Table 8-4 lists some of the default transport protocol settings.

When you select Device as the deployment method, deployment is affected if you configure a transport server for the device, such as an AUS or Configuration Engine. When using an intermediate transport server, configuration deployment goes through the server. For more information on using an intermediate server, see Deploying to a Device through an Intermediate Server, page 8-10.

Deployment can also be affected if you made out-of-band changes to the device since the last deployment. For more information, see Understanding How Out-of-Band Changes are Handled, page 8-12.

During deployment, Security Manager sends only the changes made since the last deployment to the device.

Caution You must configure at least one policy on a device before deploying to that device. If you deploy to a device without assigning at least one policy, the device's current configuration is overwritten with a blank configuration.

Table 8-4 Default Deployment Transport Protocols

Device Type	Transport Protocol	Description
ASA, IOS 12.3 and higher routers, FWSM, PIX Firewall, IPS sensors	SSL (HTTPS) (Default)	Security Manager deploys the configuration to the device using the Secure Socket Layer (SSL) protocol, otherwise known as HTTPS. With this protocol, Security Manager encrypts the configuration file and sends it to the device.

Table 8-4 Default Deployment Transport Protocols (continued)

Device Type	Transport Protocol	Description
Catalyst 6500/7600 and other Catalyst switches	SSH	Security Manager deploys the configuration to the device using a Secure Shell (SSH). This provides strong authentication and secure communications over insecure channels. Security Manager supports both SSHv1.5 and SSHv2. Once connected to the device, Security Manager determines which version to use and downloads using that version.
IOS 12.2 and 12.1 routers	Telnet	Security Manager deploys the configuration to the device using the Telnet protocol.

Related Topics

- Managing Device Communication Settings and Certificates, page 9-4
- Handling Device OS Version Mismatches, page 8-13

Deploying to a Device through an Intermediate Server

Deploying configurations through an intermediate server, such as an Auto Update Server (AUS), Cisco Networking Services (CNS) Configuration Engine, or Token Management Server (TMS), is a version of deploying directly to device. When selecting the deployment method, select Device. Security Manager sends the configuration updates to the intermediate server, where the device retrieves it (for AUS and CNS), or where you can download it to an eToken (for TMS).

You must use an intermediate server if you are using dynamic IP addresses for your device interfaces (that is, the IP addresses are provided by a DHCP server). You can also use them with static IP addresses. However, you cannot use Configuration Engine to manage IOS devices with dynamic IP addresses if you configure features that use interactive CLI commands. The following features are affected:

- Certificate Enrollment:
  - **crypto pki trustpoint**
  - **crypto isakmp client configuration group**
  - **crypto key generate rsa**
- IPS signature configuration (**ip ips signature-category**)
- IP Authproxy Banner (**ip auth-proxy-banner**)
- Catalyst device interface switchport (**interface switchport**)

Security Manager uses an intermediate server if you have configured the device to use one. The following topics describe the required configuration steps when using an intermediate server:

- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43

Deployment can be affected if you made out-of-band changes to the device since the last deployment. For more information, see Understanding How Out-of-Band Changes are Handled, page 8-12.

During deployment, Security Manager sends configuration changes based on the type of server:

- Auto Update Server (standalone or running on Configuration Engine) for PIX and ASA devices—Security Manager sends the full configuration to Auto Update Server, where the device retrieves it. The delta configuration is not sent.
- Configuration Engine for IOS devices—Security Manager sends the delta configuration to Configuration Engine, where the device retrieves it.
- TMS—Security Manager sends the delta configuration to the TMS server, from which it can be downloaded to an eToken to be loaded onto the device.

Related Topics

- Managing Device Communication Settings and Certificates, page 9-4
- Device Communication Page, page 11-21

Deploying to a File

If you choose to deploy configurations to configuration files, Security Manager creates two files: *device_name* _delta.cfg for the delta configuration, and *device_name* _full.cfg for the full configuration. If the files are created by a job that was generated from a deployment schedule, the name includes a time stamp. Configuration files are in TFTP format so that you can upload them to your devices using TFTP.

Tip You cannot deploy configurations to file for IPS devices.

If you deploy to file, you are responsible for transferring the configurations to your devices. Security Manager assumes that you have done this, so the next time you deploy to the same devices, the generated incremental commands are based on the configurations from the previous deployment. If for some reason the last change was not applied to the device, the new delta configuration will not bring the device configuration up to the one reflected in Security Manager.

Caution Although Security Manager in one sense assumes that you applied the delta configuration, in another sense, it assumes that it cannot know if the delta was deployed. Thus, Security Manager maintains an internal view of the configuration based on the last deployment made directly to the device. So, when you apply the delta, those delta changes will be considered out-of-band changes. On next deployment to the device, your out-of-band change setting might cancel the deployment. If you mix deployments to file with deployments to device, you should rediscover policies after applying file deployments to the device. For more information, see Understanding How Out-of-Band Changes are Handled, page 8-12.

To set a default directory for file deployments, select **Tools > Security Manager Administration**, then select **Deployment** (see Deployment Page, page 11-13). If you select File for the default deployment method, you also select the default directory. When you create a deployment job, you can change this directory for that job.

Deploying configurations to a file is useful when the devices are not yet in place in your network (known as green field deployment), if you have your own mechanisms in place to transfer configurations to your devices, or if you want to delay deployment. When deploying to a file, the deployment job might fail if you select a large number of devices or several devices that have large configuration files. If you encounter deployment failures, resubmit the job with fewer devices selected.

Tip Do not use commands that require interaction with the device during deployment when deploying to file. We recommend previewing your configuration before deployment to make sure there are no such commands in the file. For more information, see Previewing Configurations, page 8-44.

Understanding How Out-of-Band Changes are Handled

Security Manager considers an out-of-band change to be any change made to a device manually or outside of Security Manager control, for example, by logging into the device directly and entering configuration commands through the CLI. Paradoxically, this includes the application of delta changes that Security Manager creates when you deploy configurations to file rather than to the device.

If you are deploying to the device (rather than to file), and the deploy to device method is configured to compare the new configuration to the current configuration on the device, you can specify how to handle out-of-band changes when they are detected using the **Out of Band Change Behavior** setting. The setting does not apply when deploying to file.

This setting is ignored if you are comparing the new device configuration with the latest version stored in the Security Manager Configuration Archive. The default way to handle out-of-band changes, is set in Tools > Security Manager Administration > Deployment; for more information see Deployment Page, page 11-13. Look for the **Deploy to Device Reference Configuration** and **When Out of Band Changes Detected** settings.

Tip When the deployment method is configured to use the reference configuration in Configuration Archive, out-of-band changes are always removed. This is equivalent to selecting **Do not check for changes**.

Your options for handling out-of-band changes are:

- **Overwrite changes and show warning**—When configurations are deployed, Security Manager uploads the device's current configuration and compares it against the configuration it has in its database. If changes were made to the device manually, Security Manager continues with the deployment and displays a warning notifying you of this action. Out-of-band changes are removed from the device.
- **Cancel deployment**—When configurations are deployed, Security Manager uploads the device's current configuration and compares it against the configuration it has in its database. If changes were made to the device manually, Security Manager cancels the deployment and displays a warning notifying you of this action. You must either manually remove the out-of-band changes, or configure the same settings in Security Manager, before you can deploy configuration changes to the device.
- **Do not check for changes**—Security Manager does not check for changes and deploys the changes to the device. No warnings are issued, and any out-of-band changes are removed from the device configuration.

Before you deploy configurations, you might want to detect whether there are out of band changes on a device and analyze whether you want to recreate those changes in Security Manager policies, or allow Security Manager to overwrite the changes. For more information, see Detecting and Analyzing Out of Band Changes, page 8-45.

Related Topics

- Deploying Directly to a Device, page 8-9
- Deploying to a Device through an Intermediate Server, page 8-10
- Deploying to a File, page 8-11

Handling Device OS Version Mismatches

Before deploying a changed configuration file directly to a device, Security Manager normally uploads the current running configuration file from the device and checks the OS version running on the device with the OS version stored in the Security Manager database (you can configure it so that the archived configuration is used instead of the configuration from the device). Security Manager takes action depending on whether the OS versions match or differ from each other.

In some cases, Security Manager deploys the configuration and issues a warning, but in other cases, Security Manager cannot deploy the configuration. Security Manager deploys the configuration when:

- The device has a newer minor version, for example, ASA 8.1(2) instead of the 8.1(1), indicated in Security Manager.
- The device has a down-level minor version, for example, ASA 8.1(1) instead of 8.1(2).

Security Manager does not deploy the configuration when the device is running a new major version of the OS (for example, ASA 8.0 instead of the 7.2 indicated in Security Manager) or if the device is running a down-level major version (7.2 instead of 8.0).

The following table lists the possible actions Security Manager takes depending on the whether the OS versions match or differ from each other. The table uses the ASA device as an example; however, the actions apply to all supported device types.

Table 8-5 Deployment Action Based on OS Version Match or Mismatch

Scenario	OS Version in Security Manager Database	OS Version On Device	OS Version Used In Deployment	Action
Versions match	ASA 8.2(1)	ASA 8.2(1)	ASA 8.2(1)	Deployment proceeds with no warnings.
Device has newer minor OS version.	ASA 8.1(1)	ASA 8.1(2)	ASA 8.1(2)	Security Manager warns that it has detected a different OS version on the device than the one in the Security Manager database. Security Manager generates the CLI based on the OS version running on the device.

Table 8-5 Deployment Action Based on OS Version Match or Mismatch (continued)

Scenario	OS Version in Security Manager Database	OS Version On Device	OS Version Used In Deployment	Action
Device has newer minor OS version, one that is not directly supported by Security Manager.	ASA 8.0(2)	ASA 8.0(4)	ASA 8.0(3)	Security Manager warns that it has detected a different OS version on the device than the one in the Security Manager database. Security Manager generates the CLI based on the OS version that it supports to which the running OS version is downward-compatible.
Device has a new major OS version.	ASA 7.2(4)	ASA 8.2(1)	None. Deployment fails.	Security Manager reports an error indicating that it has detected a different OS version on the device than the one in the Security Manager database. Security Manager cannot proceed until you correct this mismatch. Remove the device from the inventory, add it again, and discover the device policies.

Table 8-5 Deployment Action Based on OS Version Match or Mismatch (continued)

Scenario	OS Version in Security Manager Database	OS Version On Device	OS Version Used In Deployment	Action
Device has an older minor OS version.	ASA 8.1(2)	ASA 8.1(1)	ASA 8.1(1)	Security Manager warns that it has detected a different OS version on the device than the one in the Security Manager database. Security Manager generates the CLI based on the OS version running on the device.
Device has an older major OS version	ASA 8.2(1)	ASA 7.2(4)	None. Deployment fails.	Security Manager reports an error indicating that it has detected a different OS version on the device than the one in the Security Manager database. Security Manager cannot proceed until you correct this mismatch. Remove the device from the inventory, add it again, and discover the device policies.

Overview of the Deployment Manager and Configuration Archive

The Deployment Manager and Configuration Archive are the main tools that you can use to manage deployment and device configurations. The following topics provide an overview of these tools:

- Understanding What You Can Do with the Deployment Manager, page 8-16
- Deployment Manager Window, page 8-16
- Deployment Schedules Tab, Deployment Manager, page 8-21
- Configuration Archive Window, page 8-23

Understanding What You Can Do with the Deployment Manager

The Deployment Manager, where you create and manage deployment jobs and schedules, provides the following benefits:

- Previewing and comparing configurations—Before you deploy a configuration file to a device, you can preview the proposed configuration file. You can also compare the proposed configuration file to what was last imported from the device or what is currently running on the device.

  After successful deployment to a device, you can view a transcript of the configuration commands downloaded and the device's responses. For more information, see Previewing Configurations, page 8-44.

- Aborting deployment jobs—You can stop a deployment job even if it is currently running. However, aborting a job that is in process does not roll back the configuration on devices that have already been reconfigured, or on devices that are in the process of being reconfigured. Only devices for which deployment has not started are prevented from being reconfigured. For more information, see Aborting Deployment Jobs, page 8-55.

- Rolling back to a previous configuration—If you deploy configurations to devices, and then determine that there is something wrong with the new configurations, you can revert to and deploy the previous configurations for those devices. For more information, see Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69.

- Viewing deployment job status—You can display information about the deployment to specific devices, including information about errors, the proposed configuration, and the transcript of the download. For more information, see Viewing Deployment Status and History for Jobs and Schedules, page 8-26.

- Scheduling deployment jobs—You can create deployment schedules to spawn deployment jobs at regular intervals. In Workflow mode, you can also schedule a deployment job to start at a future time when you deploy the job. Scheduling jobs lets you plan deployments for times when traffic on devices is low. For more information, see these topics:

  - Creating or Editing Deployment Schedules, page 8-55
  - Deploying Configurations in Workflow Mode, page 8-34

- Logging deployment job history (Workflow mode only)—You can view the history of transactions for a job. The transactions show the changes in job status initiated by various users, such as job approval, and the comments related to those status changes. For more information, see Viewing Deployment Status and History for Jobs and Schedules, page 8-26.

Deployment Manager Window

Use the Deployment Manager window to manage deployment jobs and schedules. You can display a list of deployment jobs, view job details, deploy and redeploy configurations to devices, abort deployment jobs, roll back to previous configurations on selected devices, and create schedules to automatically generate deployment jobs. You can also track changes made to deployment jobs and schedules.

Note The buttons available in the Deployment Manager depend on the Workflow mode you are using.

Navigation Path

Click the **Deployment Manager** button on the Main toolbar or select **Manage > Deployments**.

Related Topics

- Overview of the Deployment Process, page 8-1
- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Managing Device Communication Settings and Certificates, page 9-4

Field Reference

Table 8-6 Deployment Manager Window (Workflow Mode)

Element	Description
Deployment Jobs Tab	
This tab shows individual deployment jobs. Select a job in the upper pane to view its details in the tabs in the lower pane.	
Filter Options	
Beginning with 4.14, Cisco Security Manager provides filter options to search deployment jobs based on Name (deployment job name), Status, Changed By, and Device Name. After specifying the filter criteria, click Apply. The grid displays the search result. Select a job in the table to view its details in the tabs in the lower pane.	
Name	The name of the job.
Last Action	The date and time that the job or status was changed based on the time zone of the server, not the time zone of the client.
Status	The state of each job. The possible states differ based on workflow mode. For a description of the states, see the following topics: • Job States in Non-Workflow Mode, page 8-4 • Job States in Workflow Mode, page 8-6
Changed By	The name of the user who modified the job.
Description	The description of the job. Double-click the icon to see the description in a separate dialog box.
Job Type	The type of job with respect to scheduling. A **one time** job was not created from a regularly recurring job, whereas a **recurring** job was.
Create button (Workflow mode only.)	In Workflow mode, click this button to create a new job. The Create a Job dialog box opens. See Creating and Editing Deployment Jobs, page 8-35.
Open button (Workflow mode only.)	In Workflow mode, click this button to open the selected job. The Edit a Job dialog box opens. See Creating and Editing Deployment Jobs, page 8-35.

Table 8-6 Deployment Manager Window (Workflow Mode) (continued)

Element	Description
Close button (Workflow mode only.)	In Workflow mode, click this button to close and save all changes made while the selected job was open. You can close a job when it is in the Edit Open or the Submit Open state. Normally, you do not need to close a job, because you will typically submit, approve, deploy, or schedule the job for deployment. However, if the Security Manager server is suddenly unavailable or your login session times out, a job might be left in the Edit Open state. If this happens, you can close it manually by selecting it and clicking Close.
Submit button (Workflow mode only.)	In Workflow mode, click this button to submit the selected job for approval. You can submit a job when it is in the Edit or the Edit Open state. The Submit Deployment Job dialog box opens. See Submitting Deployment Jobs, page 8-38. This button is active only if you are using Workflow mode with a deployment job approver.
Reject button (Workflow mode only.)	In Workflow mode, click this button to reject the selected job if you are not satisfied with the configurations generated for the devices. You can reject jobs only in workflow mode with a deployment job approver. After a job is rejected, it can be opened for editing or discarded. See Approving and Rejecting Deployment Jobs, page 8-39. You are prompted to enter an optional comment to explain why you are rejecting the job.
Approve button (Workflow mode only.)	In Workflow mode, click this button to approve the selected job. After a job is approved, it can be deployed. See Approving and Rejecting Deployment Jobs, page 8-39. You are prompted to enter an optional comment to explain why you are approving the job.
Discard button (Workflow mode only.)	In Workflow mode, click this button to discard the selected job. You can discard a job when it is in any state except Deployed, Deployment Failed, or Aborted. Once discarded, the job cannot be edited, submitted, approved, or deployed. The job state is shown as discarded until the job is purged from the system either automatically as set on the Workflow settings page or manually (for more information, see Workflow Page, page 11-75). You are prompted to enter an optional comment to explain why you are discarding the job. See Discarding Deployment Jobs, page 8-41.

Table 8-6 Deployment Manager Window (Workflow Mode) (continued)

Element	Description
Deploy button (All modes.)	Click this button to deploy generated CLI commands devices or files. The behavior of this button differs depending on Workflow mode: • (Non-Workflow mode.) Click this button to create a deployment job. If you have unsubmitted changes, you are first prompted to submit them. The Deploy Saved Changes dialog box opens, where you can select which devices to include in the job. Note that this button does not act on the deployment job selected in the table, if any; instead, it creates a new deployment job. See Deploying Configurations in Non-Workflow Mode, page 8-28. • (Workflow mode.) Click this button to deploy the selected job. If the job is in the Approved state, the Deploy Job dialog box opens (see Deploying a Deployment Job in Workflow Mode, page 8-39). If the job is in the deployed, failed, or aborted state then the Redeploy Job dialog box opens. See Redeploying Configurations to Devices, page 8-53.
Generate Report button (All modes.)	Click this button to create a deployment status report for the selected job. You can generate the report in HTML and PDF formats. Jobs must be in deployed, failed, rolled back, or aborted state. The deployment status report includes a summary of the job plus the full and delta configurations and the job transcript. You can use this report for your own purposes or to aid in troubleshooting a problem with Cisco TAC. For more information, see Generating Deployment or Discovery Status Reports, page 10-30.
Refresh button (All modes.)	Click this button to reload job information from the Security Manager server. If the message *Auto Refresh is On* is displayed beneath the table, the job list is automatically refreshed periodically. **Note** The auto refresh setting is configured in the administration settings for deployment: select **Tools > Security Manager Administration > Deployment**.
Redeploy button (Non-Workflow mode only.)	In Non-Workflow mode, click this button to redeploy the selected job, which deploys the same generated CLI commands to the same devices or files selected in the original job. The Redeploy Job dialog box opens. See Redeploying Configurations to Devices, page 8-53. (In Workflow mode, click the Deploy button to redeploy configurations for the selected job.)
Abort button (All modes.)	Click this button to abort the selected job if it is in the Deploying, Scheduled, or Rolling Back state. A warning asks you to confirm the action. See Aborting Deployment Jobs, page 8-55.
Rollback button (All modes.)	Click this button to deploy the previously deployed configuration to the devices in the selected job. The Deployment Rollback dialog box opens (see Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69).

Table 8-6 Deployment Manager Window (Workflow Mode) (continued)

Element	Description
Summary tab	Displays summary information about the status of the selected deployment job, such as the status of the job, the name of the deployment job, the number of devices included in the job, the number of devices deployed successfully, and the number of devices deployed with errors.
Details tab	Displays detailed information for the selected job. The table lists each device included in the job, whether deployment succeeded or failed, the tickets containing changes that are part of the job for the device, and a summary of the number of warnings, errors, or failures for the device. Select a device in the table to view the results for that device: • Double-click the icon in the Config column to view the configuration (see Previewing Configurations, page 8-44). If you deleted the device from the inventory, the configuration and transcript might not be available. • If you were deploying to the device, double-click the icon in the Transcripts column to view a transcript of the commands sent to the device and the device's responses. See Viewing Deployment Transcripts, page 8-62. • If Ticket Management is enabled, the Last Ticket(s) column displays the ticket IDs of the tickets containing changes that are part of the deployment for the device. You can click on the Ticket IDs to view additional information about the ticket, such as Creator and Last Modified date. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72). • When you select a device, the Messages box in the lower left contains a summary of the messages generated for the deployment. Select an item to view its description to the right. You might have to enlarge the window to make the Description box visible. If applicable, there might also be information on the actions you can take to resolve the problems.
History tab (Workflow mode only.)	Displays a log of the changes that have been made to the selected job. The information includes the state changes, the user who made the change, the date and time of the change (based on the Security Manager server time), and any comments the user entered to document the change.
Deployment Schedules Tab Use this tab to schedule regular deployment jobs. For detailed information about this tab, see Deployment Schedules Tab, Deployment Manager, page 8-21.	

Deployment Workflow Commentary Dialog Box

When you perform an action in the Deployment Manager while working in Workflow mode, you are prompted to enter a comment to describe the action. The comments are preserved in the history for the job or schedule.

The title of the dialog box indicates the action you are taking. Enter an optional comment and click **OK** to perform the action.

Navigation Path

In Workflow mode, select a job or schedule in the Deployment Manager and click the appropriate button to perform the desired action.

Deployment Schedules Tab, Deployment Manager

Use the Deployment Schedules tab on the Deployment Manager window to create regularly recurring deployment jobs. Whenever the scheduled deployment time occurs, Security Manager creates a specific deployment job based on the scheduled job.

Navigation Path

Click the **Deployment Manager** button on the Main toolbar or select **Manage > Deployments**, and then click the **Deployment Schedules** tab in the upper pane.

Related Topics

- Overview of the Deployment Process, page 8-1
- Creating or Editing Deployment Schedules, page 8-55
- Suspending or Resuming Deployment Schedules, page 8-58

Field Reference

Table 8-7 Deployment Schedules Tab, Deployment Manager Window

Element	Description
Deployment Schedule Table	
This table shows deployment job schedules. Select a schedule in the table to view its details in the tabs in the lower pane.	
Filter Options	
Beginning with 4.14, Cisco Security Manager provides filter options to search deployment schedules based on Name (deployment schedule name), Status, and Device Name. After specifying the filter criteria, click Apply. The grid displays the search result. Select a schedule in the table to view its details in the tabs in the lower pane.	
Name	Name of the job schedule. Jobs created from this schedule use this name plus a time stamp.
Status	The status of the schedule: • Edit—In Workflow mode, the schedule is being created. You can open it and change its settings. No jobs are created from schedules that are being edited. • Active—Deployment jobs will be created according to this schedule. • Suspended—The schedule was suspended and no jobs are being created by it. You can restart the schedule by selecting it and clicking **Resume**.

Table 8-7 Deployment Schedules Tab, Deployment Manager Window (continued)

Element	Description
Recurrence	How often deployment jobs will be created from this schedule.
Next Run	The date and time a deployment job will next be created from this schedule.
Last Run	The date and time of the most recent deployment job created from this schedule.
Schedule End	The date and time the schedule is no longer active. If the schedule has no end date, Active Indefinitely is indicated.
Description	The description of the job schedule. Double-click the icon to see the description.
Create button	Click this button to create a deployment job schedule. The Schedule dialog box opens where you can create the schedule (see Schedule Dialog Box, page 8-56.
Open button	Click this button to open the selected schedule. The Schedule dialog box opens where you can view or modify the schedule (see Schedule Dialog Box, page 8-56). In non-Workflow mode, modifying the schedule does not change its status. In Workflow mode, the status changes to Edit, and you must resubmit it for approval.
Close button (Workflow mode only)	Click this button to close and save all changes made while the schedule was open. You can close a schedule when it is in the Edit Open or the Submit Open state. Typically, you will have to close schedules only if the Security Manager server becomes unavailable while you have a schedule open.
Submit button (Workflow mode only)	Click this button to submit the selected schedule for approval if you are operating in Workflow mode with an approver. You can submit a schedule when it is in the Edit or the Edit Open state. You are prompted for an optional comment to explain the submission, and an e-mail is generated to the approver in Workflow mode.
Reject button (Workflow mode only)	Click this button to reject the selected schedule. You are prompted for an optional comment to explain the rejection, and an e-mail is generated to the approver and submitter in Workflow mode.
Approve button (Workflow mode only)	Click this button to approve the selected schedule. You are prompted for an optional comment to explain the approval, and an e-mail is generated to the approver and submitter in Workflow mode.
Discard button	Click this button to discard the selected schedule. You can discard a schedule unless there is an active deployment job that was created from the schedule. (You can wait for the job to finish, or abort the job and then discard the schedule.) You are prompted for an optional comment to explain the discard, and an e-mail is generated to the approver and submitter in Workflow mode.

Table 8-7 Deployment Schedules Tab, Deployment Manager Window (continued)

Element	Description
Refresh button	Click this button to reload schedule information from the Security Manager server. If the message *Auto Refresh is On* is displayed beneath the table, the schedule list is automatically refreshed periodically. **Note** The auto refresh setting is configured in the administration settings for deployment: select **Tools > Security Manager Administration > Deployment**.
Suspend button	Click this button to suspend the selected schedule. Suspending the schedule does not delete the schedule, but it prevents the creation of deployment jobs based on it. You are prompted for a comment to explain the suspension, and an e-mail is generated to the approver in Workflow mode.
Resume button	Click this button to reactivate a suspended schedule. You are prompted for a comment to explain the suspension, and an e-mail is generated to the approver in Workflow mode.
Summary tab	Displays summary information about the selected schedule. Besides the fields shown in the table, summary information includes the number of devices included in the schedule and the user ID of the person who last changed the schedule.
Devices tab	Displays the devices that are included in the selected schedule. These are the devices to which configurations are deployed when a deployment job is created from the schedule. To change the device list, click **Open**, then click **Add Devices** on the Schedule dialog box.
History tab	Displays a log of the changes that have been made to the selected schedule. The information includes the state changes, the user who made the change, the date and time of the change (based on the Security Manager server time), and any comments the user entered to document the change.
Jobs tab	Displays a list of the deployment jobs that have been created based on the selected schedule. Information includes the name of the job, the date and time the job was created based on server time (not the client time), and the job status. If you select a job, and click the **Deployment Job** tab, the selected job is highlighted and you can view the job details. For information on job status, see these topics: • Job States in Workflow Mode, page 8-6 • Job States in Non-Workflow Mode, page 8-4

Configuration Archive Window

The Configuration Archive stores configuration versions for each device managed by Security Manager. If you delete a device from Security Manager, all of the device's configurations are also deleted from the Configuration Archive.

You can use Configuration Archive to:

- View the transcript of a configuration deployment for a selected device.

- View and compare configuration versions.
- View CLI differences between deployed configuration versions.
- Roll back to an earlier configuration version, provided that the configuration originated from the device. You should roll back configurations only under extreme circumstances. For more information, see these topics:
  - Understanding Configuration Rollback, page 8-63
  - Using Rollback to Deploy Archived Configurations, page 8-70
- Add the current running configuration for a device to the archive.

You can sort the list of configuration versions for a device by clicking on the column heading that you want to sort on. Clicking the column heading toggles between sorting the rows in ascending or descending order. You can also control the fields displayed by right-clicking on any column heading and selecting or deselecting the desired column names under the Show Columns command.

Navigation Path

Select **Manage > Configuration Archive**.

Related Topics

- Configuration Archive Page, page 11-6
- Viewing and Comparing Archived Configuration Versions, page 8-59
- Understanding Configuration Rollback, page 8-63
- Using Rollback to Deploy Archived Configurations, page 8-70
- Understanding Rollback for Devices in Multiple Context Mode, page 8-64
- Understanding Rollback for Failover Devices, page 8-65
- Understanding Rollback for Catalyst 6500/7600 Devices, page 8-65
- Understanding Rollback for IPS and IOS IPS, page 8-66
- Adding Configuration Versions from a Device to the Configuration Archive, page 8-59
- Filtering Items in Selectors, page 1-45

Field Reference

Table 8-8 Configuration Archive Window

Element	Description
Device Selector	Lists the devices in the device inventory. Select a device to see the configuration versions for the device that are available in the archive. These are displayed in the right pane.
Version ID	The version number of the configuration version. By default, this column is not displayed. To display it, right click any column heading and select **Show Columns > Version ID**.
Created On	The date and time that the configuration version was archived.

Table 8-8 Configuration Archive Window (continued)

Element	Description
Created By	The user ID or system ID associated with adding the configuration version to the archive. If there are two names in the form *username1*(*username2*), the first name is the user who initiated the request, and the name in parentheses is the system identity user. For more information on the system identity trust user, see the *Installation Guide for Cisco Security Manager*.
Archival Source	The origin of the archiving event (for example, User Request, Deployment or Provision, Discovery).
Creation Comment	A description about how or why the configuration version was created.
Transcript Icon	When double-clicked, displays a transcript of a configuration version that was deployed to a device. A transcript is the log file of transactions between Security Manager and a device captured during a deployment or rollback operation. It includes commands sent and received between server and device from the time of the deployment or rollback request, but it does not include communication that occurs during the initial discovery phase of deployment, when Security Manager obtains the current configuration from the device.
View button	Click this button to display the selected configuration in the Config Version Viewer window (see Configuration Version Viewer, page 8-60), where you can also compare the configuration to other configuration versions.
Rollback button	Click this button to roll the device configuration back to the selected configuration version, provided that the configuration originated from the device. You should roll back configurations only under extreme circumstances. For more information see these topics: • Understanding Configuration Rollback, page 8-63 • Using Rollback to Deploy Archived Configurations, page 8-70
Add from Device button	Click this button to have Security Manager retrieve the current running configuration from the device and add it as a configuration version to the archive. This is useful for any device whose configuration might have been changed directly in its CLI. For more information on adding configuration versions, see Adding Configuration Versions from a Device to the Configuration Archive, page 8-59.

Working with Deployment and the Configuration Archive

The following topics provide information about managing deployment and using the Configuration Archive:

- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Tips for Successful Deployment Jobs, page 8-27

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Previewing Configurations, page 8-44
- Detecting and Analyzing Out of Band Changes, page 8-45
- Redeploying Configurations to Devices, page 8-53
- Aborting Deployment Jobs, page 8-55
- Creating or Editing Deployment Schedules, page 8-55
- Suspending or Resuming Deployment Schedules, page 8-58
- Adding Configuration Versions from a Device to the Configuration Archive, page 8-59
- Viewing and Comparing Archived Configuration Versions, page 8-59
- Viewing Deployment Transcripts, page 8-62

Viewing Deployment Status and History for Jobs and Schedules

Using the Deployment Manager, you can view status and history information for deployment jobs and schedules, as well as create and manage them. To open the Deployment Manager window, select **Manage > Deployments**.

Jobs and schedules are displayed on separate tabs. However, as jobs are created based on a deployment schedule, those jobs appear in the regular jobs list. Click the appropriate tab to view the list of jobs or schedules, where you can see this information:

- Deployment Jobs—The top pane displays a list of the deployment jobs. If you select a job, more detailed information appears in the lower pane:
  - Summary tab—The Summary tab shows information such as the job status, number of devices deployed successfully, and number of devices deployed with errors.
  - Details tab—The Details tab shows the status details for each device in the deployment.
  - History tab (Workflow mode only)—The History tab displays transactions that occurred to the selected job since it was created. Each row in the table shows the action that occurred, the user who performed the action, the date and time it occurred, and comments, if any, that the user entered.
- Deployment Schedules—The top pane displays a list of the deployment schedules. If you select a schedule, more detailed information appears in the lower pane:
  - Summary tab—The Summary tab shows information such as the schedule, the time of the next job to be created from the schedule, the time a job was last run based on the schedule, the number of devices included in the schedule and the user ID of the person who last changed the schedule.
  - Devices tab—The Devices tab shows the list of devices that are included in the schedule.
  - History tab—The History tab shows the state changes and related comments of the schedule. You can track which user performed each action.
  - Jobs tab—The Jobs tab shows a list of deployment jobs that were created from the schedule and their statuses. You can also view these jobs on the Deployment Jobs tab.

The status information in the Deployment Manager window refreshes automatically unless you turned off automatic refresh in the Security Manager Administration Deployment page (Tools > Security Manager Administration > Deployment). A message below the job or schedule table indicates whether automatic refresh is on. If it is off, refresh status information by clicking **Refresh**.

Related Topics

- Overview of the Deployment Process, page 8-1
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Previewing Configurations, page 8-44
- Redeploying Configurations to Devices, page 8-53
- Aborting Deployment Jobs, page 8-55
- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Creating or Editing Deployment Schedules, page 8-55
- Suspending or Resuming Deployment Schedules, page 8-58

Tips for Successful Deployment Jobs

Successful deployment depends on many things, as explained in Troubleshooting Deployment, page 9-9. In addition to factors involving network communications and the proper functioning of the device, you can also improve the results of deployment by keeping the following tips in mind when you select devices for a deployment job or start the job:

- You must configure at least one policy on a device before deploying to that device. If you deploy to a device without assigning at least one policy, the device's current configuration is overwritten with a blank configuration and the device will be non-functional.
- Firewall devices only—If you manually added a firewall device (as described in Adding Devices by Manual Definition, page 3-26), we highly recommend that you discover (import) the factory-default policies for that device before deploying to that device. Bringing these policies into Security Manager prevents you from unintentionally removing them the first time you deploy to that device. For more information about factory-default policies for firewall devices, see Default Firewall Configurations, page 46-2. For more information about importing policies, see Discovering Policies, page 5-12.
- Deployment might take from a few minutes to an hour or more, depending on the number of devices in the deployment job.
- Modifying a subset of devices that are part of a VPN might make the VPN inoperable. If you select a subset of devices that are part of a VPN when creating a deployment job, you are warned and given the opportunity to select the other devices in the VPN. See Warning - Partial VPN Deployment Dialog Box, page 8-32.
- You cannot select devices that were included in other deployment jobs that are in an active state (Edit, Edit Open, and Approved). You can select devices that were included in other deployment jobs that are in the Deployed, Failed, Discarded, or Aborted states.

- Firewall service modules (FWSMs) and Intrusion Detection System service modules (IDSMs) contain virtual devices. Security Manager considers the module and the virtual devices to be separate devices.
- Some changes to the FWSM might require the Catalyst Multiservice function card (MSFC) to be updated as well. If you select an FWSM that has these types of changes, Security Manager notifies you that you must include the MSFC in the deployment job, and it will select the MSFC device for you automatically. However, if the MSFC is already included in another active deployment job, you cannot include the MSFC in the current deployment job. You must remove the MSFC from the other deployment job, discard the other deployment job, or include the FWSM in the other deployment job.
- The status of deployments to Catalyst 6500/7600 devices shows deployment to the device as well as its interface contexts when policy changes contain interface commands that affect the interface contexts (child devices). This can occur when you deploy a policy change that affects a VLAN in which the switch participates or when you update inventory, for example, by adding or deleting interface contexts.

Related Topics

- Overview of the Deployment Process, page 8-1
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Creating and Editing Deployment Jobs, page 8-35
- Managing Device Communication Settings and Certificates, page 9-4
- Detecting and Analyzing Out of Band Changes, page 8-45
- Job States in Non-Workflow Mode, page 8-4
- Job States in Workflow Mode, page 8-6

Deploying Configurations in Non-Workflow Mode

When you deploy configurations, you can transfer them to devices either directly or to another transport server (such as AUS, CNS, or TMS) in the network or create them as configuration files in a directory on the Security Manager server. See Understanding Deployment Methods, page 8-8 for more information.

Note If you have made changes to a Unified ACL entry through Policy Object Manager which is used in a RAVPN policy (DAP, Group Policy, and alike) on a device, the device and the ticket does not get displayed in the Deploy Saved Changes window. You must click Add other devices and add the device manually.

Tip Before creating the deployment job, read Tips for Successful Deployment Jobs, page 8-27. That topic includes tips and cautions you should keep in mind when creating deployment jobs.

Caution You must configure at least one policy on a device before deploying to that device. If you deploy to a device without assigning at least one policy, the device's current configuration is overwritten with a blank configuration.

Note When using virtual sensors: An IPS device and all of the virtual sensors on it must be deployed as a group. If you make changes to a virtual sensor and then deploy it, Security Manager deploys the parent device and all virtual sensors associated with it.

Before You Begin

- Make sure that devices have been bootstrapped. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are deploying to a transport server, such as AUS, CNS, or TMS, make sure the server, Security Manager settings, and device have been set up properly.

Related Topics

- Overview of the Deployment Process, page 8-1
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Managing Device Communication Settings and Certificates, page 9-4
- Understanding How Out-of-Band Changes are Handled, page 8-12

Step 1 Do one of the following in non-Workflow mode:

- Select **File > Submit and Deploy** or click the **Submit and Deploy Changes** button on the toolbar.

  Note These options are not available when Ticket Management is enabled.

- Select **File > Deploy**.
- Click the **Deployment Manager** button on the Main toolbar and click the **Deployment Jobs** tab if it is not active. Click **Deploy**.

Security Manager validates all of the policy changes that were made since the last deployment. If the validation results in errors, resolve the errors before attempting to deploy again. If there are only warnings or informational messages, click **OK** to proceed to the Deploy Saved Changes dialog box.

Step 2 In the Deploy Saved Changes dialog box, do the following:

- Select the devices to which you want to deploy configurations. The device selector lists all devices for which policy changes were made but not yet deployed, and initially all changed devices are selected for deployment.

  All device groups that contain changed devices are shown, and you can select or deselect the devices using the device group folder. If you select or deselect a device that appears in more than one group, it is selected or deselected in all groups; however, a device is deployed to only once in the job. Right-click and select **Expand All** to open all of the folders.

  The Deploy Saved Changes dialog box shows the date, time, and user associated with the changes that will be included in the deployment for the selected devices. This information changes based on the devices you select for deployment. If you have Ticket Management enabled, the tickets

associated with the changes to be deployed are also displayed. You can click a ticket ID to view details of the ticket and to navigate to an external ticket management system if configured (see Ticket Management Page, page 11-72).

If you detect out of band changes, when you close the OOB Changes dialog box, the device names are color-coded based on the results: green indicates out of band change; red indicates an error during the detection process; no color change indicates no out of band changes.

- If you want to add devices that do not have policy changes to the deployment job, click **Add other devices** to open the Add Other Devices dialog box (see Add Other Devices Dialog Box, page 8-58). You might want to add unchanged devices if a device was manually modified and you want to return the device to its previous configuration (the one stored in the Security Manager database).
- (Optional) To change the method used to deploy configurations, click **Edit Deploy Method** to open the Edit Deployment Method dialog box (see Edit Deploy Method Dialog Box, page 8-30). There is a system default deployment method (which your organization chooses), so you might not need to change the method. You can select these methods:
  - Device—Deploys the configuration directly to the device or to the transport mechanism specified for the device. For more information, see Deploying Directly to a Device, page 8-9 or Deploying to a Device through an Intermediate Server, page 8-10.
  - File—Deploys the configuration file to a directory you select on the Security Manager server. For more information, see Deploying to a File, page 8-11.

Before proceeding with the deployment, you can do the following:

- Preview proposed configurations and compare them against last deployed configurations or current running configurations. Right-click the device and select **Preview Config**. For more information, see Previewing Configurations, page 8-44.
- Analyze the devices for out of band changes by clicking the **Detect OOB Changes** button. For more information, see Detecting and Analyzing Out of Band Changes, page 8-45 and OOB (Out of Band) Changes Dialog Box, page 8-48.

Step 3 Click **Deploy** to start the deployment job for the selected devices, which generates the required configuration files and applies them according to your selected deployment method.

The Deployment Status Details dialog box opens so that you can view the status of the deployment. It displays summary information about the job, status about the deployment to each device, and messages indicating why the deployment failed.

In the Deployment Details table, select a row corresponding to a device to display deployment status messages specifically for that device. For more information, see Deployment Status Details Dialog Box, page 8-32.

If deployment to any device failed, you can redeploy configurations to the failed devices. For more information, see Redeploying Configurations to Devices, page 8-53.

Edit Deploy Method Dialog Box

Use the Edit Deploy Method dialog box to specify whether to deploy the generated configurations directly to the devices in the network or to create configuration files in a directory on the Security Manager server.

Navigation Path

Click **Edit Deploy Method** in the Deployment—Create or Edit a Job dialog box (Workflow mode) or Deploy Saved Changes dialog box (non-Workflow mode). For the procedures, see:

- Creating and Editing Deployment Jobs, page 8-35
- Deploying Configurations in Non-Workflow Mode, page 8-28

Related Topics

- Understanding Deployment Methods, page 8-8
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Managing Device Communication Settings and Certificates, page 9-4

Field Reference

Table 8-9 Edit Deploy Method Dialog Box

Element	Description
Device	The name of the device.
Method	The deployment method to use: • Device—Deploys the configuration directly to the device or to the transport mechanism specified for the device. For more information, see Deploying Directly to a Device, page 8-9 or Deploying to a Device through an Intermediate Server, page 8-10. • File—Deploys the configuration file to a directory on the Security Manager server. If you select File, specify the directory to which you want to deploy the configuration file in the Destination column. You cannot use file deployment with IPS devices. For more information, see Deploying to a File, page 8-11. **Note** To set the deployment method for more than one device at a time, select the desired rows, right-click and select **Edit Selected Deploy Method**. The Edit Selected Deploy Method dialog box opens where you can make your selections.
Destination	If you selected File in the Method field, enter the directory to which you want to deploy the configuration file. Click **Browse** to select from a list of available directories.
Preview Config button	Click this button to display the proposed configuration changes for the selected device. You can compare it to the last deployed configuration or the current running configuration. For more information, see Previewing Configurations, page 8-44.
Out of Band Change Behavior	Click the radio button corresponding to the action you want Security Manager to take regarding changes made directly on the device using the CLI. For a complete explanation of how to handle out-of-band changes, including the meaning of the available options, see Understanding How Out-of-Band Changes are Handled, page 8-12.

Warning - Partial VPN Deployment Dialog Box

Use the Partial VPN Deployment dialog box to select other devices that are part of a VPN to which you are deploying configurations.

When you create a deployment job and the job contains devices in a VPN, you must select all of the devices in the VPN. If you select a subset of devices and try to deploy to only those devices, this dialog box appears so that you can select the other devices that are part of the VPN.

Navigation Path

- Non-Workflow mode—If you select a subset of devices in a VPN in the Deploy Saved Changes dialog box, this dialog box appears when you click **Deploy**.
- Workflow mode—If you select a subset of devices in a VPN in the Create or Edit a Job dialog box, this dialog box appears when you click **OK**.

Related Topics

- Creating and Editing Deployment Jobs, page 8-35
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Field Reference

Table 8-10 Partial VPN Deployment Warning Dialog Box

Element	Description
VPN	The name of the VPN.
Missing Devices	All the devices in the VPN that were not selected for deployment.
Is Device in Other Job	Whether the missing device is part of another deployment job.
Deploy to All Devices in VPN button	Click this button to deploy to all devices in the VPN. You can deploy to all devices in the VPN only if the devices are not in other deployment jobs.
Deploy to Selected Devices button	Click this button to deploy only to the devices selected in the Create or Edit a Job or Deploy Saved Changes dialog boxes.

Deployment Status Details Dialog Box

The Deployment Status Details dialog box appears while configurations are being deployed to selected devices. It displays summary information about the job, status about the deployment to each device, and messages indicating why the deployment failed.

In the Deployment Details table, select a row corresponding to a device to display deployment status messages for that device.

Note You can click **Close** to close this dialog box and continue working in Security Manager while deployment continues.

Navigation Path

From the Deploy Saved Changes dialog box, click **Deploy**.

Related Topics

- Overview of the Deployment Process, page 8-1
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Tips for Successful Deployment Jobs, page 8-27
- Managing Device Communication Settings and Certificates, page 9-4
- Device Communication Page, page 11-21

Field Reference

Table 8-11 Deployment Status Details Dialog Box

Element	Description
Deployment Status Details	
Progress Status Bar	A visual representation and percentage of devices that were successfully updated.
Status	The status of the deployment. The possible states are Deploying, Aborted, Successful, and Failed. For descriptions of these states, see Job States in Non-Workflow Mode, page 8-4.
Deployment Job Name	The name of the deployment job.
Devices To Be Deployed	The total number of devices in the deployment job.
Devices Deployed Successfully	The number of devices that were updated successfully.
Devices Deployed With Errors	The number of devices that failed to be updated.
Deployment Details This table lists the devices that are included in the deployment job.	
Device	The name of the device.
Status	The status of the deployment to the device. For descriptions of these states, see Job States in Non-Workflow Mode, page 8-4.
Summary	The number of warnings, errors, and failures for the device.
Method	The method of deployment to the device. Possible methods are File and Device.
Config	The device configuration file. Double click the icon to preview the configuration for a device. For more information, see Previewing Configurations, page 8-44.
Transcript	The commands Security Manager issued to the device and the responses from the device during deployment if you are deploying to the device (instead of deploying to a file). Double-click the icon to see the transcript for a device.
Last Ticket(s)	The tickets containing changes that are part of the deployment for the device. You can click on the Ticket IDs to view additional information about the ticket, such as Creator and Last Modified date. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).

Table 8-11 Deployment Status Details Dialog Box (continued)

Element	Description
Messages	The warning, error, and failure messages, as indicated by the severity icon. When you select an item, the Description box to the right describes the message in detail. The Action box to the right provides information on how you can correct the problem.
Generate Report button	Click this button to create a deployment status report for this job. You can generate the report in HTML and PDF formats. The report includes a summary of the job plus the full and delta configurations and the job transcript. You can use this report for your own purposes or to aid in troubleshooting a problem with Cisco TAC. For more information, see Generating Deployment or Discovery Status Reports, page 10-30.
Refresh button	Click this button to update the status information.
Abort button	Click this button to abort the deployment job. You can abort deployment jobs only while they are in the Deploying, Scheduled, or Rolling Back state. Aborting a job stops deployment of configuration files to pending devices, but has no effect on devices to which deployments are in progress (commands are being written to a device) or to which deployment has already completed successfully.

Deploying Configurations in Workflow Mode

The task of deploying configurations in Workflow mode is a multiple step process. You must create a deployment job, get it approved, and then deploy the job. This process ensures that organizations that separate task authorizations among personnel can implement their control processes.

When you deploy configurations, you can transfer them to devices either directly or to another transport server (such as AUS, CNS, or TMS) in the network or create them as configuration files in a directory on the Security Manager server. See Understanding Deployment Methods, page 8-8 for more information.

Tip Before creating the deployment job, read Tips for Successful Deployment Jobs, page 8-27. That topic includes tips and cautions you should keep in mind when creating deployment jobs.

Before You Begin

- Make sure that devices have been bootstrapped. For more information, see Chapter 2, “Preparing Devices for Management”.
- If you are deploying to a transport server, such as AUS, CNS, or TMS, make sure the server, Security Manager settings, and device have been set up properly.

Related Topics

- Overview of the Deployment Process, page 8-1
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41

- Deploying Configurations to a Token Management Server, page 8-43
- Managing Device Communication Settings and Certificates, page 9-4
- Understanding How Out-of-Band Changes are Handled, page 8-12

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Create the deployment job. Click **Create** and enter the job properties. For the procedure, see Creating and Editing Deployment Jobs, page 8-35.

When you finish creating a job, you can select whether to submit it. If you are not using a deployment job approver, you can also automatically submit, approve, and deploy the job, in which case you do not need to complete the other steps in this procedure.

Step 3 (Workflow with approver) Submit the job. If you did not submit the job, select it in the Deployment Manager window and click **Submit**. An e-mail is sent to the approver. For more information, see Submitting Deployment Jobs, page 8-38.

Step 4 (Workflow with or without an approver) Approve the job. If you did not approve the job when you created it, select it in the Deployment Manager window and click **Approve**. If there is a separate person who approves jobs, that person must perform this step. For more information, see Approving and Rejecting Deployment Jobs, page 8-39.

Step 5 (Workflow with or without an approver) Deploy the job. If you did not deploy the job when you created it, select it in the Deployment Manager window and click **Deploy**. You can specify a future time to start the job, or start it immediately, and configurations are deployed according to the properties of the job. For more information, see Deploying a Deployment Job in Workflow Mode, page 8-39

Note You can discard a deployment job at any time before you deploy it. For more information, see Discarding Deployment Jobs, page 8-41.

Creating and Editing Deployment Jobs

In Workflow mode, before you deploy policy configurations to your devices, you must create a deployment job. When you create a job, you select the devices to which you want to deploy the configurations, whether you want to deploy directly to the devices or to an output file, and when you want the job to take place.

Note If you have made changes to a Unified ACL entry through Policy Object Manager which is used in a RAVPN policy (DAP, Group Policy, and alike) on a device, the device and the ticket does not get displayed in the Deployment- Create a Job window. You must click Add other devices and add the device manually.

Tip Before creating the deployment job, read Tips for Successful Deployment Jobs, page 8-27. That topic includes tips and cautions you should keep in mind when creating deployment jobs.

Caution You must configure at least one policy on a device before deploying to that device. If you deploy to a device without assigning at least one policy, the device's current configuration is overwritten with a blank configuration.

Before You Begin

- Make sure that devices have been bootstrapped. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are deploying to a transport server, such as AUS, CNS, or TMS, make sure the server, Security Manager settings, and device have been set up properly.

Related Topics

- Overview of the Deployment Process, page 8-1
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Understanding How Out-of-Band Changes are Handled, page 8-12
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Do one of the following:

- Click **Create** to create a new job.
- Select an editable job and click **Open** to edit the job. You cannot edit a job that has already been deployed.

The Create a Job or Edit a Job dialog box opens.

Step 3 In the dialog box, do the following to define the contents of the job:

- **Job Name and Description**—Keep the default job name or enter a more meaningful name. Because the job name enables you to distinguish one job from another, you should assign a name that reflects the contents of the job. You cannot change the name after you create the job. Optionally, enter a description of the job.
- Select the devices to which you want to deploy configurations. The device selector lists all devices for which policy changes were made but not yet deployed, and initially all changed devices are selected for deployment.

  All device groups that contain changed devices are shown, and you can select or deselect the devices using the device group folder. If you select or deselect a device that appears in more than one group, it is selected or deselected in all groups; however, a device is deployed to only once in the job. Right-click and select **Expand All** to open all of the folders.

  If you detect out of band changes, when you close the OOB Changes dialog box, the device names are color-coded based on the results: green indicates out of band change; red indicates an error during the detection process; no color change indicates no out of band changes.
- If you want to add devices that do not have policy changes to the deployment job, click **Add other devices** to open the Add Other Devices dialog box (see Add Other Devices Dialog Box, page 8-58). You might want to add unchanged devices if a device was manually modified and you want to return the device to its previous configuration (the one stored in the Security Manager database).

- (Optional) To change the method used to deploy configurations, click **Edit Deploy Method** to open the Edit Deployment Method dialog box (see Edit Deploy Method Dialog Box, page 8-30). There is a system default deployment method (which your organization chooses), so you might not need to change the method. You can select these methods:
  - Device—Deploys the configuration directly to the device or to the transport mechanism specified for the device. For more information, see Deploying Directly to a Device, page 8-9 or Deploying to a Device through an Intermediate Server, page 8-10.
  - File—Deploys the configuration file to a directory you select on the Security Manager server. For more information, see Deploying to a File, page 8-11.

Before proceeding with the deployment, you can do the following:

- Preview proposed configurations and compare them against last deployed configurations or current running configurations. Right-click the device and select **Preview Config**. For more information, see Previewing Configurations, page 8-44.
- Analyze the devices for out of band changes by clicking the **Detect OOB Changes** button. For more information, see Detecting and Analyzing Out of Band Changes, page 8-45 and OOB (Out of Band) Changes Dialog Box, page 8-48.

Step 4 Select how you want the job handled when you close the dialog box. The options available to you depend on whether you are using Workflow mode with a deployment job approver:

- **Without an approver**—If you are not using a separate approver, you have these options:
  - **Close the job**—Close the job and leave it in the edit state. Select this option if you know you want to make additional modifications to the job.
  - **Approve the job**—Close the job and approve it but do not deploy it. Configure the following:
  - **Comments**—(Optional) Comments about the job approval.
  - **Submitter**—The e-mail address of the person submitting the job for approval. Notifications of job state changes are sent to this address, which is initially the e-mail address associated with the user account you used to log into Security Manager. Ensure that the address is the correct one so that you receive the notifications.
  - **Deploy the job**—Close the job, approve it, and deploy it. Configure the following:
  - **Options**—Whether to Deploy Now or Schedule. If you select Schedule, additional fields appear where you can specify the date and time when the job should be run. The time is in 24-hour format and is based on the time zone of the Security Manager server, which is not necessarily the same time zone that you are currently in. The target time must be at least five minutes in the future.
  - **Comments**—(Optional) Comments about the deployment job.
  - **Send Deployment Status Notification**—Whether Security Manager should send e-mail notifications whenever the job status changes.

    If you select this option, enter the e-mail addresses of the people who should receive notifications in the Job Completion Recipients field. If you enter multiple addresses, separate them with commas. The field initially contains the default approver and your e-mail addresses.
- **With an approver**—If you are using a separate approver, you can configure the following options:
  - **Submit the job**—Whether to submit the job for approval. By default this check box is selected.
  - **Approver E-mail**—The e-mail address of the approver if you are submitting the job for approval. The default approver e-mail address is entered in the field, but you can change it.
  - **Comments**—(Optional) Comments you want to send to the approver, if any.

– **Submitter E-mail**—The e-mail address of the submitter. The field initially contains the e-mail address associated with the user account you used to log in, but you can change it to another address.

Step 5 Click **OK**.

Depending on your selection for how to handle the job, you might still need to submit, approve, and deploy the job. See these topics for more information:

- Submitting Deployment Jobs, page 8-38
- Approving and Rejecting Deployment Jobs, page 8-39
- Deploying a Deployment Job in Workflow Mode, page 8-39

Submitting Deployment Jobs

In some organizations, before jobs can be deployed, they must be approved by a separate user with the appropriate permissions. In this case, Workflow mode is enabled *with* a deployment job approver, and you must submit the job to this user for review. The user reviews the job and either approves or rejects it.

If you are using Workflow mode *without* a deployment job approver, you can review and approve the job yourself. You do not submit jobs in this mode. For more information, see Approving and Rejecting Deployment Jobs, page 8-39.

Note You enable and disable deployment job approval under Tools > Security Manager Administration > Workflow. For more information, see Workflow Page, page 11-75.

This procedure assumes that you already created the job. You can also submit the job when you create it by selecting the **Submit the job** checkbox in the Create a Job dialog box.

Related Topics

- Deployment Manager Window, page 8-16
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Status window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the job to submit.

Step 3 Click **Submit**.

The Submit Deployment Job dialog box opens.

Step 4 Enter the following information:

- **Approver**—The e-mail address of the person to be notified of your submission. The default approver e-mail address is entered in the field, but you can change it.
- **Comment**—(Optional) Comments you want to send to the approver, if any.
- **Submitter**—The e-mail address of the person submitting the deployment job. The field initially contains the e-mail address associated with the username you used to log into Security Manager, but you can change it to another e-mail address.

Step 5 Click **OK**.

The job status changes to Submitted. The approver must approve the job before you can deploy it.

Approving and Rejecting Deployment Jobs

In some organizations, before jobs can be deployed, they must be approved by a separate user with the appropriate permissions. In Workflow mode *with* a deployment job approver, one user submits a job, and another one previews the job and either approves or rejects it.

In Workflow mode without a deployment job approver, you can create and approve the job at the same time. For more information, see Creating and Editing Deployment Jobs, page 8-35.

When you reject a job, the devices in the job immediately become available for inclusion in other jobs. A rejected job cannot be deployed, but it can be opened for viewing and editing.

Note You enable and disable deployment job approval under Tools > Security Manager Administration > Workflow. For more information, see Workflow Page, page 11-75

Related Topics

- Deployment Manager Window, page 8-16
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select a submitted job and do one of the following:

- Click **Approve**.
- Click **Reject**.

You are prompted for an optional comment for your action. The comments are preserved in the history for the job. After submitting your comment, an e-mail notification is sent (if e-mail notifications are configured) and the job status changes to Approved or Rejected, as appropriate. The job can now be deployed (see Deploying a Deployment Job in Workflow Mode, page 8-39).

Deploying a Deployment Job in Workflow Mode

When you work in Workflow mode, to deploy configurations to devices you must create a deployment job and have it approved. If you are working without a separate approver, you can approve and deploy the job yourself. Otherwise, you must submit it to an approver.

Deploying a deployment job in workflow mode simply starts a job. You cannot change the contents of a job during deployment.

Note Deployment might take from a few minutes to an hour or more, depending on the number of devices in the deployment job.

Before You Begin

- Make sure that devices have been bootstrapped. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are deploying to a transport server, such as AUS, CNS, or TMS, make sure the server, Security Manager settings, and device have been set up properly.
- Create a job. For more information, see Creating and Editing Deployment Jobs, page 8-35.
- If using Workflow mode with a deployment job approver, submit the job. For more information, see Submitting Deployment Jobs, page 8-38.
- Approve the job. For more information, see Approving and Rejecting Deployment Jobs, page 8-39.

Related Topics

- Overview of the Deployment Process, page 8-1
- Deployment Manager Window, page 8-16
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Managing Device Communication Settings and Certificates, page 9-4

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the job to deploy.

Step 3 Click **Deploy**.

The Deploy Job dialog box opens.

Step 4 In the Deploy Job dialog box, make these selections:

- **Options**—How you want to run the job. Select **Schedule** to run the job at some point in the future. Select **Deploy Now** to run the job immediately. If you schedule the job for a future time, the changes deployed in the job are based on the changes that existed when the job was created, not when the job is run.

  If you select Schedule, date and time fields appear.

  - Click the calendar icon to pick the day on which to run the job.
  - In the Time field, enter the time to start the job in 24-hour clock format. The time must be in time zone of the Security Manager server, which is not necessarily the same as your time zone. The time must be at least 5 minutes in the future.
- **Comments**—(Optional) An explanation of why you are deploying the job.
- **Require Deployment Status Notifications, Job Completion Recipients**— Whether Security Manager should send an e-mail when the job status changes.

  If you elect to send status notifications, enter the recipient's e-mail address. The field initially contains the e-mail address associated with the user account you used to log in. You can enter multiple addresses by separating them with commas.

Step 5 Click **OK**.

You are returned to the Deployment Manager window. The job status changes to Deploying. When the deployment is complete, the job status changes to Deployed.

Discarding Deployment Jobs

In Workflow mode, you can discard a job when it is in any state except Deployed, Deployment Failed, or Aborted. The job state is shown as discarded until the job is purged from the system, either automatically as set on the Workflow Management page or manually.

Related Topics

- Deployment Manager Window, page 8-16
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the job to discard.

Step 3 Click **Discard**. You are prompted for an optional comment to explain why you are discarding the job.

Deploying Configurations Using an Auto Update Server or CNS Configuration Engine

If your organization uses Auto Update Server (AUS) or Cisco Networking Services (CNS) Configuration Engine to manage the deployment of configurations to your network devices, you can use these intermediate servers with Security Manager. To perform this type of deployment, you need to set up the device, the AUS or Configuration Engine, and Security Manager properly. This procedure explains the tasks that you need to perform.

Tip You cannot successfully deploy a configuration to AUS that requires Security Manager to download other files to the device. For example, some remote access VPN policies allow you to configure plug-ins, Anyconnect clients, and Cisco Secure Desktop configurations. These files are not sent to AUS. Do not use AUS if you want to configure these types of policy.

Related Topics

- Overview of the Deployment Process, page 8-1
- Chapter 2, “Preparing Devices for Management”
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Managing Device Communication Settings and Certificates, page 9-4

Step 1 Set up the AUS or Configuration Engine using the documentation for those products.

Step 2 Configure the devices to use the server. The following topics describe the configuration steps depending on the type of server and the desired setup:

- Setting Up AUS on PIX Firewall and ASA Devices, page 2-8

Step 3 When you add the device to Security Manager, select the AUS or Configuration Engine for the device if your chosen method allows it. If the AUS or Configuration Engine is not already defined in Security Manager, you can identify it to Security Manager as you add the network device. For detailed procedures, see these topics:

- Adding Devices from the Network, page 3-12
- Adding Devices from Configuration Files, page 3-22
- Adding Devices by Manual Definition, page 3-26
- Adding Devices from an Inventory File, page 3-31
- Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36

Tip After you add a device to the Security Manager inventory, you can change the assigned server in the device properties. Right-click the device and select **Device Properties**. Configure the server using the device properties if you could not identify it while adding the device.

Step 4 For devices that are using AUS, configure the AUS policy for the device in Security Manager. Do one of the following:

- Configure the policy for a single device. In Device view, select the device, and then select **Platform > Device Admin > Server Access > AUS** from the Device Policy selector.
- Configure a shared policy that you can assign to many devices that share the same AUS. In Policy view, select **PIX/ASA/FWSM Platform > Device Admin > Server Access > AUS** from the Policy Types selector. Right-click **AUS** and select **New AUS Policy** to create a policy, or select an existing policy from the Policies selector to change the policy. Select the Assignments tab to assign the policy to specific devices.

The server you identify in this policy must also be the server you identify in the device properties. The device properties identify the server to which Security Manager will send the configuration, whereas the AUS policy defines the server the device will contact.

Tip If you change AUS servers, keep in mind that the device will continue to use the AUS server defined in its current configuration until it receives a new configuration. Thus, you should change the AUS policy but deploy the configuration using the previous AUS server. After deployment is successful, change the device properties to point to the new server.

Step 5 In Security Manager, deploy your configurations using the **Deploy to Device** deployment method. Security Manager sends the configuration to the AUS or Configuration Engine, where the network device retrieves it.

Depending on the Workflow mode you are using, follow these procedures:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Deploying Configurations to a Token Management Server

If your organization requires the use of a Token Management Server (TMS) for applying configuration updates to routers, you can use Security Manager in conjunction with your TMS processes. To perform this type of deployment, you need to set up the device, TMS, and Security Manager properly. This procedure explains the tasks that you need to perform.

Related Topics

- Overview of the Deployment Process, page 8-1
- Chapter 2, "Preparing Devices for Management"
- Including Devices in Deployment Jobs or Schedules, page 8-8
- Understanding Deployment Methods, page 8-8
- Managing Device Communication Settings and Certificates, page 9-4

Step 1 Set up the TMS as an FTP server. Security Manager uses FTP to deploy the configuration file to the TMS, from which it can be downloaded and encrypted onto an eToken. The eToken can then be connected to the USB port of a router and the configuration downloaded. See the TMS product documentation for more information.

Step 2 In Security Manager, select **Tools > Security Manager Administration > Token Management** to identify the TMS server to Security Manager.

By default, Security Manager uses the Security Manager server as the TMS, but you can specify a different server. You must enter the hostname or IP address, a username and password for the TMS, the directory to which configuration files should be copied, and the public key file location in Security Manager. For more information, see Token Management Page, page 11-73.

Step 3 Specify TMS as the transport protocol to be used for Cisco IOS routers.

You can set this parameter globally for all Cisco IOS routers or for a specific device:

- Globally—Select **Tools > Security Manager Administration > Device Communication** and select TMS in **Transport Protocol (IOS Routers 12.3 and above)**.
- Device—Right click the device in the Device selector and select **Device Properties**. On the General tab, select TMS as the transport protocol in the Device Communications Group. Because not all routers support TMS, you might not be able to configure TMS for specific devices.

Step 4 In Security Manager, deploy your configurations using the **Deploy to Device** deployment method. Security Manager sends the delta configuration to the TMS server.

Depending on the Workflow mode you are using, follow these procedures:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Step 5 Using the TMS, download the configuration to the eToken. See the TMS product documentation for more information.

Step 6 Download the configuration from the eToken to the router and save the configuration to the device. Plug the eToken into the router, then enter the following commands to download the configuration to the router, where *usb_token_id* is either **usbtoken0** or **usbtoken1**, depending on which USB port you used. The default PIN is 1234567890.

```
router# crypto pki token usb_token_id login PIN
router# config terminal
```

```
router(config)# crypto pki token default secondary config CCCD
router(config)# exit
router# write memory
```

Tip CCCD is the private sector on the eToken where the configuration file resides. When you enter the **crypto pki token default secondary config CCCD** command, the CLI on the e-token merges with the CLI on the router.

Previewing Configurations

There are many ways to preview a device configuration. You can select a device from the Device selector and select **Tools > Preview Configuration**, or you can click the **Preview Config** button in several dialog boxes.

Tip You can also right click a device in Map view and select **Preview Configuration**.

When you preview a configuration, the configuration is displayed in the Config Version Viewer dialog box. The proposed configuration is on the left. You can select to view the delta configuration (which shows the changes since the last deployment) or the full configuration. You can also compare the configuration to the last one deployed to the device or the current running configuration in the right pane.

The contents of the proposed configuration can differ depending on where you view it from:

- If you use **Tools > Preview Configuration**, or right click the device in the Device selector and select **Preview Configuration**, the proposed configuration includes changes that you have not yet submitted to the database.
- If you preview the configuration while creating a deployment job, the proposed configuration includes only those changes that you have submitted to the database. These are the changes that will be deployed to the device if you start the deployment job.

The value of previewing configurations is that it lets you see the actual device commands that will be used to configure the device. If you are a CLI expert, this can help you verify that you are getting the configuration you expect. Even if you are not a CLI expert, you can use the information to look for more information in the command reference for the operating system on Cisco.com.

Note New objects are created when you rediscover ASA 8.2.3 devices through Cisco Security Manager after configuring network objects on the device. Ignore the CLIs for these new objects that are generated in preview configuration.

Following are some tips for previewing configurations:

- Many configuration options are specific to one or more interfaces. If you must specify an interface name in a policy, the previewed configuration will include the commands for the policy only if the Interfaces policy defines the interface that you specify. Ensure that you configure the Interfaces policy before previewing configurations.
- If you preview the configuration for a virtual sensor, the preview that you see is for the parent device, not the virtual sensor, because the configuration for a virtual sensor is stored on the parent device.

- If you are just curious about what commands a policy will configure, consider adding a dummy device and configure the policy for that device. This will help prevent unintended configuration changes in your real devices. To add a dummy device, use the procedure described in Adding Devices by Manual Definition, page 3-26.
- Policies are validated before you can view the configuration. The validation happens for all devices, not just the device you are previewing. Thus, you might see errors and warnings that apply to different devices. If any errors or warnings occur, the Preview Messages dialog box appears. The dialog box lists all of the messages, including their severity and possible solutions. Click **OK** to continue to the Config Version Viewer dialog box. Click **Details** to see detailed information on the problems.

The following table explains the fields in the Config Version Viewer window used for previewing configurations.

Table 8-12 Config Version Viewer (Preview Configuration) Dialog Box

Element	Description
Proposed Config Type	The type of configuration you want to view. For example, you can view the full configuration or just the delta (the changes from the last deployed configuration). The proposed configuration is displayed in the left pane.
Compare to Version	Choose a configuration to compare against the proposed configuration. The selected configuration is displayed in the right pane. • None—Leaves the reference configuration blank. • Last Deployed—Displays the last configuration that was deployed to the device and compares it with the proposed configuration. • Running Config—Displays the current configuration running on the device and compares it with the proposed configuration. The device must be accessible to obtain the running configuration.
First Difference button	Moves the cursor to the first difference noted between the proposed and reference configurations.
Previous Difference button	Moves the cursor to the previous difference noted between the proposed and reference configurations.
Current Difference button	Centers the currently selected difference on the page.
Next Difference button	Moves the cursor to the next difference noted between the proposed and reference configurations.
Last Difference button	Moves the cursor to the last difference noted between the proposed and reference configurations.
Print button	Prints the configuration.

Detecting and Analyzing Out of Band Changes

When you deploy configurations to devices, Security Manager either removes out of band changes or cancels the deployment based on your deployment settings. (For a detailed explanation of out of band changes and how they are handled during deployment, see Understanding How Out-of-Band Changes are Handled, page 8-12.)

In the most typical scenario, Security Manager removes any out of band changes during deployment. However, there might have been good reason for those changes to have been made to the device outside of Security Manager. Thus, it is good practice to analyze a device for out of band changes before deploying configurations, so that you have the opportunity to proactively recreate any configuration changes that should be preserved.

There are many ways to detect whether there have been out of band changes to a device configuration since the last deployment to the device. If there have been changes (by another device management application or by direct CLI updates), you can preview the changes and determine whether to update the device policies before deployment, or to deploy and overwrite those out of band changes. (Out of band changes are sometimes called OOB changes in Security Manager). In some scenarios, the OOB changes are not detected. For information on handling such exceptions, see Exceptions to Out of Band Change Detection, page 8-47.

Tip Out of band change detection is available only for IOS, ASA, PIX, FWSM devices, and security contexts; it is not available for IPS devices. However, the settings for handling out of band changes during deployment also apply to IPS devices; the difference is that you cannot proactively analyze these changes in IPS devices prior to deployment.

To determine whether there have been out of band changes on one or more device, do any of the following in Device view:

- Select **Tools > Detect Out of Band Changes**. You are prompted to select the devices to evaluate for out of band changes. Select the devices or device groups, click **>>** to move them to the selected list, and click **OK**. For more information on selecting devices, see Using Selectors, page 1-45.
- Select one or more devices or device groups, right-click and select **Detect Out of Band Changes**. The selected devices are evaluated for changes.
- During deployment, select the devices to include in deployment and click the **Detect OOB Changes** button. (The button is available on the Deploy Saved Changes dialog box and the Deployment—Create or Edit a Job dialog box, depending on the workflow mode you are using.) The selected devices are evaluated for changes.

  For information on the deployment procedure, see:

  - Deploying Configurations in Non-Workflow Mode, page 8-28
  - Creating and Editing Deployment Jobs, page 8-35

When you start the detection process, the OOB (Out of Band) Changes Dialog Box is opened so that you can view the results. Each selected device is evaluated by retrieving the current running configuration and comparing it to the most recent configuration stored in Configuration Archive. Security Manager does not consider any unmanaged policy types when evaluating differences between the configurations.

Tip Note that if you are in the process of deployment, the running configuration is **not** compared to the one you are proposing to deploy, so if you detect out of band changes, you might also want to preview the proposed configuration to see if you already implemented the same change in Security Manager policies. Right-click the device in the deployment dialog box and select **Preview Config**. You can compare the proposed configuration to the current running configuration. For more information, see Previewing Configurations, page 8-44.

The OOB Changes dialog box shows the results of change detection. If a device has out of band changes, the icon for the device in the device selector changes to green. Select a device in the left pane of the OOB Detail tab to view the changes from the latest configuration in configuration archive. Use the buttons at the bottom of the window to move from change to change. The legend at the bottom explains the color coding used to describe the changes.

When evaluating changes, consider the following:

- If you want to keep the change, update the relevant policy in Security Manager to recreate the policy. Use preview config to ensure that your policy changes produce the desired results. Security Manager might use different naming conventions, so consider whether the policy results in the same thing, rather than being exactly the same text. Keep in mind that out of band change detection looks for syntactic differences, not semantic differences.
- If you use another application to configure certain types of policies, consider unmanaging that policy type in Security Manager. Security Manager ignores any configuration commands related to policies that it is not managing. For more information, see Policy Management Page, page 11-64.

Tip If you are detecting changes during deployment, when you close the OOB Changes dialog box, the device names in the deployment dialog box are color-coded based on the results: green indicates out of band change; red indicates an error during the detection process; no color change indicates no out of band changes.

Exceptions to Out of Band Change Detection

If you have not approved the activity in which the changes have been done, the Cisco Security Manager database will not get updated. This results in a discrepancy between the Cisco Security Manager configuration archive (that the OOB feature uses) and the Cisco Security Manager database. When you do not approve the activity, Cisco Security Manager does not detect Our of Band (OOB) changes that are applied to the device. As a result, Cisco Security Manager does not stop deployment (overwriting OOB changes) even when you have configured it to cancel deployment when OOB changes are detected (see Understanding How Out-of-Band Changes are Handled, page 8-12). This section discusses this exception and how to handle it.

The following tasks are executed in Cisco Security Manager (Workflow mode), when a policy rediscovery is initiated:

Step 1 After rediscovery (Discovering Policies on Devices Already in Security Manager, page 5-15), a new device configuration is written to the Cisco Security Manager configuration archive.

Step 2 If you do not approve the activity in which policy rediscovery is done, the Cisco Security Manager database will not be updated with the new device configuration and will continue to use old configuration data. Thus, there will be a mismatch between the configuration archive and the Cisco Security Manager database. This might result in the OOB changes on the device to be overwritten by Cisco Security Manager. This occurs even when you have configured it to cancel deployment when OOB changes are detected (see Understanding How Out-of-Band Changes are Handled, page 8-12).

Note If the policy rediscovery activity is not approved, Out of Band (OOB) changes are not detected between the Cisco Security Manager database and configuration on the device. This is because OOB changes are detected using the Cisco Security Manager configuration archive, which has been updated with discovered configuration from device (Step 1 above). On the other hand, Cisco Security Manager database still has the previous configuration of device, since activity has not been approved.

In addition, when you are Previewing Configurations for the discovered device before the activity is approved, the preview configuration does not show correct configuration changes. In order to see correct differences, you must approve the activity first or preview configurations from another activity.

To overcome these exceptions, do the following:

Step 1 Create a new activity for rediscovery.

Step 2 On completion of policy rediscovery, Submit and Approve the activity. Verify if the activity has been approved.

Step 3 To confirm if configuration changes for the rediscovered device is displayed as expected, perform Previewing Configurations for the device.

Step 4 Deploy changes from Cisco Security Manager to the device, if required.

OOB (Out of Band) Changes Dialog Box

Use the OOB Changes dialog box to view and analyze out of band changes on a device. An out of band change is any difference between the running configuration on a device and the most recent configuration for the device stored in Configuration Archive. Note that Security Manager considers only managed policy types when evaluating whether there is a difference between the configurations.

Tip Configurations are compared for syntactic differences, not semantic differences. Thus, functionally equivalent configurations might be identified as out of band changes.

Tip As an example, consider a simple case in which configuration lines are swapped in the device configuration without making any changes to the semantics. In this simple example case: 1) if there is an object group in Security Manager at line number 100, and 2) the same object group exists in the ASA configuration at any line number *except* 100, then 3) Security Manager detects and reports the change as OOB. To summarize this simple example case, Security Manager reports an OOB change even though this change in the order of a few configuration lines did not result in any changes to the semantics.

There are two tabs on this dialog box:

- **OOB Detail**—This tab shows the detailed results and progress of the detection process. The fields are described below.

- **OOB Summary**—This tab shows a summary of the detection results, and becomes available only after the detection process is completed on all selected devices. The information is by device and includes a time stamp (date, time, time zone) and difference data that indicates the additions, subtractions, and changes with the associated configuration line number. You can select text on this tab, use Ctrl+click to copy it to the clipboard, and paste it in another application (such as NotePad).

For more information on detecting and analyzing out of band changes, see Detecting and Analyzing Out of Band Changes, page 8-45. For more information on handling out of band changes during deployment, see Understanding How Out-of-Band Changes are Handled, page 8-12.

Beginning with Version 4.7, Security Manager has a tool to help you re-sync out of band changes. For more information on this new tool, see OOB Re-sync. Tool, page 8-50

Navigation Path

There are several ways to start the out of band change detection process. You can use the **Tools > Detect Out of Band Changes** command, or select one or more devices right-click and select **Detect Out of Band Changes**. You can also click the **Detect OOB Changes** button in Deploy Saved Changes dialog box and Deployment—Create or Edit a Job dialog box as explained in the following procedures:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Creating and Editing Deployment Jobs, page 8-35

Related Topics

- Previewing Configurations, page 8-44
- Filtering Items in Selectors, page 1-45

Field Reference

Table 8-13 OOB Changes Dialog Box

Element	Description
Selected Devices list (left pane)	This list contains all devices you selected to evaluate for out of band changes, organized in device groups (if any). Select a device to see the results in the right pane. The icons for the devices change color based on the results of the detection process: • Green—There are out of band changes. • Red—The out of band detection process failed for some reason. • No color change—No out of band changes.
Configuration Comparison (right pane)	The right pane shows the results of the change detection process for the selected device. Messages will indicate if OOB detection is still in progress, if there are no changes, or if there was an error that prevented change detection from completing. If there are changes, the right pane shows both the running configuration retrieved from the device and the latest configuration for the device stored in Configuration Archive. The legend at the bottom of the window describes the color coding used to indicate changes, and you can use the following buttons to move from change to change.
First Difference button	Moves the cursor to the first difference noted between the configurations.

Table 8-13 OOB Changes Dialog Box (continued)

Element	Description
Previous Difference button	Moves the cursor to the previous difference noted between the configurations.
Current Difference button	Centers the currently selected difference on the page.
Next Difference button	Moves the cursor to the next difference noted between the configurations.
Last Difference button	Moves the cursor to the last difference noted between the configurations.

OOB Re-sync. Tool

The OOB Re-sync Tool, which is new in Security Manager 4.7, helps you re-sync, or reconcile, out of band data. The OOB Re-sync Tool is an extension of the OOB Detection Tool available in Security Manager 4.6 and earlier versions and continued into 4.7.

Tip Out of band (OOB) data is the difference between the last archived configuration of Security Manager and the latest configuration running on the device. OOB data comes into existence—with the result that you need to update the device CLI—for reasons such as these: 1) emergency requirements (primarily for ACLs) mean that there is no time to use Security Manager and complete its workflow process because an unknown validation error is blocking deployment; 2) the use of management applications other than Security Manager to manage the same devices; and 3) <100% feature support by Security Manager for ASAs means that some ASA features need to be managed by using the CLI. Any changes made to a device using a third-party tool together with any CLI changes made to a device sum up as OOB data.

The OOB Re-sync Tool aims to automate the process of bringing OOB data on a device into your Security Manager installation while retaining the policy structures that you previously established.

Without the OOB Re-sync Tool, Security Manager (versions 4.6 and earlier) has only the following administrative options when OOB data is detected during deployment:

- Warn and override the OOB changes (Default)—Cisco Security Manager detects for OOB changes during deployment, warns the user of the OOB changes but goes ahead and negates/wipes out the OOB changes.
- Stop deployment—Aborts the deployment when OOB changes are detected.
- Do not check for OOB changes—OOB changes are not even detected during deployment and are overridden on the device.

The following objects are supported by the OOB Re-sync Tool:

- Network Object/Object-Group(s)
- Security Group(s)
- Service Object Group(s)
- User Group(s)
- Time Range Object(s)

Note The OOB Re-sync Tool does not support OOB changes for routers.

The OOB Re-sync Tool does not re-sync all objects/ACLs. It re-syncs access rules and unified access rules; it does not, for example, re-sync IPv6 access rules. Note the details for policies in the following list:

- Access rules (unified) are supported
- IPv4 access rules are supported
- IPv6-only access rules are *not* supported
- Ethertype ACLs are *not* supported
- Standard ACLs are *not* supported

The OOB Re-sync Tool has a straightforward work flow:

1. Detect OOB changes by using one of the following methods to run the existing tool:
   - Configuration Manager > Tools > Detect Out Of Band Changes... .
   - Configuration Manager > [tool bar] > Detect OOB Changes icon
   - Configuration Manager > Device View > right-click a device > click **Detect Out Of Band Changes**.
   - In the Deploy Saved Changes dialog box, click **Detect OOB**.
2. If out of band changes are detected, they appear in the right pane of the OOB (Out of Band) Changes Dialog Box on the OOB Details tab. The OOB Details tab displays a report of the changes and the target rule number, shared policy, sections, affected devices, and CLI.

   Also if out of band changes are detected, the Re-Sync Summary tab in the right pane of the OOB (Out of Band) Changes Dialog Box becomes active.

   After OOB changes are detected by the existing OOB detection tool, click Evaluate; after you do, Security Manager will further analyze the differences in the configuration running on the device and the configuration available in Security Manager. After this analysis, the Re-Sync Summary Tab becomes active. On this tab, Security Manager displays other details, such as ACE, object(s) that will be added or deleted, and rule location(s).

   Note Security Manager also will annotate the policy rule table, both in Device and Policy View and for objects in the Policy Object Manager.

3. After the Re-Sync Summary Tab becomes active, you have the option of generating a report and checking to see if there is any CLI that is not supported by OOB functionality. After checking the report, you can opt to accept the changes by clicking **Accept**, and if the operation of persisting ACL or object changes in the device went through fine, you will be prompted with a "success" message.

   Note If you modify access rules that are part of a shared policy, the OOB Re-sync Tool in this particular case will annotate both the rule that has actually been modified and the rule just above it. This occurs when you 1) modify at least two access rules that are part of a shared policy; 2) run the OOB Re-sync Tool; and 3) accept the changes. In this case, the OOB Re-sync Tool reports an OOB condition for some rules in addition to the ones that you modified. It is important to understand that the rules themselves and the shared policy are not adversely affected.

4. On the OOB Detail tab in the left pane, you can request a report in .pdf format. To do this click the **Generate Report** button. Cisco recommends that you always generate this report and save it to help with troubleshooting if needed.

Tip The following example is a brief description of a scenario in which you use the OOB Re-sync Tool and find some changes in a device that are not supported by the OOB Re-sync Tool along with access-rule changes that are supported. In this scenario, you have an ASA that uses both IPv4 and IPv6. If you use **Tools > Detect Out Of Band Changes..** and find OOB changes, you will need to manually reconcile the IPv6 changes before you can elect to re-sync. the OOB changes by using the OOB Re-sync Tool.

You should be aware of several caveats when using the OOB Re-sync Tool. These are listed in the following table.

Table 8-14 Caveats

Interface Role	Access rules tied to interface roles and multiple interfaces rules will not be absorbed. However, the rules will be annotated to help the user copy the OOB rules from the OOB re-sync report and run "Import Rules" at the right rule location to absorb the OOB CLIs.
Shared Policies	OOB changes that affect shared policies will not be re-synced since doing so would affect other devices. Rules will be annotated to help the user import rules.
Objects	The OOB re-sync process always creates overrides for objects. However, if object override is selectively disabled for that object, then re-sync will not be allowed until the user enables the device overrides for that object.
Unsupported Access List	IPv6 only access-list that existed before the introduction of unified access-list is not supported for resync Ether Type Access List re-sync is not supported
OOB Access-Group CLI	OOB changes on access-group CLI cannot be absorbed. To explain this caveat further: • Under these circumstances (OOB changes on access-group CLI), you cannot choose evaluation; that is, you cannot elect to re-sync. the OOB changes. • In a case involving both 1) OOB changes that the OOB Re-sync Tool can re-sync and 2) changes in access-group CLIs, you must resolve the changes with respect to the access-group command before you can elect to re-sync the OOB changes.
Conditional Re-sync of Remarks	ACL remarks added to an access-list CLI created out-of-band will be absorbed during re-sync as part of re-sync of the rule. However, any random OOB change in an ACL remark alone will not be absorbed during re-sync

Table 8-14 Caveats (continued)

Rule Split	Rules get split during re-sync of OOB changes done within combined rules. User needs to run "Combine Rules" on the flattened rules thus re-synced to restore to the original rule if possible

Related Topics

- Detecting and Analyzing Out of Band Changes, page 8-45
- Understanding How Out-of-Band Changes are Handled, page 8-12

Redeploying Configurations to Devices

You can redeploy a deployment job if you want to. This is especially valuable for jobs in the Failed or Aborted states. You can redeploy to all devices in the job, or you can select specific devices (such as the devices to which deployment failed).

Tips on Redeploying a Configuration to a Replacement Device

If you have to replace a device, for example, due to hardware failure, you cannot simply redeploy the last deployment job from the device, because Security Manager does not know that the device is actually a new one. To deploy the old device's configuration to the new device, you have these options:

- If the new device is the exact same model and operating system version as the replaced device, you can select the old device in the device selector, right-click and select **Preview Configuration**, and copy and paste the full configuration to the new device. However, this does not migrate certificates from the old device to the new one. You must re-enroll the device or renew the certificate yourself.
- If the new device is not exactly identical to the old device, follow the procedure described in Changes That Change the Feature Set in Security Manager, page 3-54.

Before You Begin

- Make sure that devices have been bootstrapped. For more information, see Chapter 2, "Preparing Devices for Management".
- If you are deploying to a transport server, such as AUS, CNS, or TMS, make sure the server, Security Manager settings, and device have been set up properly.

Related Topics

- Overview of the Deployment Process, page 8-1
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41
- Deploying Configurations to a Token Management Server, page 8-43
- Managing Device Communication Settings and Certificates, page 9-4
- Understanding Deployment Methods, page 8-8
- Job States in Non-Workflow Mode, page 8-4
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the job that contains the devices to which you want to redeploy configurations, then do one of the following:

- In non-Workflow mode, click **Redeploy**.
- In Workflow mode, click **Deploy**.

The Redeploy a Job dialog box opens. The dialog box lists the devices in the deployment job, showing the device name, the deployment method used, the status of the previous deployment, and the name of the deployment job that updated the device.

Step 3 In the Redeploy a Job dialog box, do the following:

- **Selection column**—Select the devices to which you want to redeploy configurations by putting checkmarks in the checkboxes in the Selection column. Initially all failed devices are selected.
- **Deployment Method, Destination**—(Optional) You can change the method used to deploy configurations for individual devices. The initially selected method is the one used in the job. You can select these methods
  - **Device**—Deploys the configuration directly to the device or to the transport mechanism specified for the device. For more information, see Deploying Directly to a Device, page 8-9 or Deploying to a Device through an Intermediate Server, page 8-10.
  - **File**—Deploys the configuration file to a directory on the Security Manager server. If you select File, specify the directory to which you want to deploy the configuration file in the Destination column. Click **Browse** to select from a list of available directories. You cannot use file deployment with IPS devices. For more information, see Deploying to a File, page 8-11.

  Note To set the deployment method for more than one device at a time, select the devices, right-click and select **Edit Selected Deploy Method**. The Edit Selected Deploy Method dialog box opens where you can make your selections.

- **Out of Band Change Behavior**—(Optional) Select how you want Security Manager to respond if it detects that changes were made on the device by someone other than Security Manager (these are called out of band changes). For a complete explanation of how to handle out-of-band changes, including the meaning of the available options, see Understanding How Out-of-Band Changes are Handled, page 8-12.

Note Before proceeding with the deployment, you can preview proposed configurations and compare them against last deployed configurations or current running configurations. Highlight the row for a device and click **Preview Config**. For more information, see Previewing Configurations, page 8-44.

Step 4 Click **OK**.

Aborting Deployment Jobs

You can stop a deployment job if you do not want to deploy the configurations or you want to postpone deployment.

You can abort deployment jobs only while they are in the Deploying, Scheduled, or Rolling Back state. Aborting a job stops deployment of configurations to pending devices, but has no effect on devices to which deployments are in progress (commands are being written to a device) or to which deployment has already completed successfully.

To abort a job, do either of the following:

- Click Abort on the Deployment Status dialog box while you are viewing the running status of an active job. See Deployment Status Details Dialog Box, page 8-32.
- Select **Manage > Deployments** to open the Deployment Manager window, then select the job on the Deployment Jobs tab and click **Abort**.

The Abort the Job dialog box opens and asks you to confirm that you want to abort the job. Click **OK** to confirm.

After you abort a job, the deployment status of pending devices changes to Aborted.

To resume deployment, redeploy the job. See Redeploying Configurations to Devices, page 8-53 for more information.

Related Topics

- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Job States in Non-Workflow Mode, page 8-4
- Job States in Workflow Mode, page 8-6

Creating or Editing Deployment Schedules

You can create deployment schedules to create deployment jobs at regular intervals. Schedules can help you ensure that the selected devices get regular configuration updates.

Tip When you include a device in a schedule, the device is included in deployment jobs that are generated from the schedule only if changes have been made to the device configuration and those changes were committed to the database. Thus, you might see changes when previewing the device configuration even though the device was not included in a scheduled deployment, if you have not submitted those changes (or if they have been submitted by not yet approved, when using a separate approver in Workflow mode).

Related Topics

- Overview of the Deployment Process, page 8-1
- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Suspending or Resuming Deployment Schedules, page 8-58

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Schedules** tab if it is not active (see Deployment Schedules Tab, Deployment Manager, page 8-21).

Step 2 Do one of the following:

- If you are creating a new schedule, click **Create**.
- If you are editing an existing schedule, select it in the Deployment Schedule table and click **Open**.

The Schedule dialog box opens (see Schedule Dialog Box, page 8-56).

Step 3 Enter at least this information in the Schedule dialog box:

- The name of the schedule.
- If you are using Workflow mode with an approver, ensure that the approver e-mail address is correct. Also verify your e-mail address (in the Submitter field), and choose whether you want to get notifications whenever the status of the job changes.
- Define the first date and time the schedule should start, and select how often deployment jobs will be generated based on the schedule. Also determine whether the schedule should have an end date, after which no new jobs are created from it.
- Click **Add Devices** and select all the devices that should be included in the deployment job. Including devices does not lock them from being modified by users or included in other deployment jobs or schedules.

  If Security Manager is configured to use user-login credentials for accessing devices, your username and password are captured during schedule creation. If you change your password, you will need to recreate the schedule.

Step 4 Click **OK**. The schedule is added to the Deployment Schedule table.

Step 5 (Workflow mode only) If you are operating in Workflow mode, you must complete these additional steps:

- If you are using an approver for deployment jobs, select the schedule in the table and click **Submit** to submit the schedule to the approver. You are prompted to verify the approver's e-mail address and to enter comments to help the approver evaluate the schedule. The approver will have to approve the schedule before it becomes active.
- If you are not using an approver, select the schedule in the table and click **Approve** to approve it yourself and to activate the schedule.

Schedule Dialog Box

Use the Schedule dialog box to create a regularly recurring deployment job.

Navigation Path

Select **Manage > Deployments** to open the Deployment Manager window, click the Deployment Schedules tab in the upper pane, and do one of the following:

- Click **Create** to create a new schedule.
- Select a schedule and click **Open** to view or modify its properties.

Related Topics

- Creating or Editing Deployment Schedules, page 8-55
- Suspending or Resuming Deployment Schedules, page 8-58

Field Reference

Table 8-15 Schedule Dialog Box

Element	Description
Schedule Name Group	
This group defines the name of the job and the job's notification requirements.	
Name	The name of the job. When individual deployment jobs are created from this schedule, a time stamp is added to the job name.
Description	The description of the purpose of the job.
Approver Email (Workflow only)	The e-mail address of the person who should approve the schedule.
Comments (Workflow only)	(Optional) Information to help the approver evaluate the schedule when you save this schedule.
Submitter Email (Workflow only)	The e-mail address of the person who is submitting this schedule for approval. This field initially contains the e-mail address associated with the user account you used to log into Security Manager, but you can change it.
Require Deployment Status Notifications (Workflow only)	Whether to send e-mail messages for any change in the job status for the job schedule or any job created from it. Messages are sent to the approver and the submitter.
Recurrence Pattern Group	
The fields in this group define the job schedule.	
Start Date	The first day of the schedule. Click the calendar icon to select the date from a calendar.
Time (Start)	The time of day to run the schedule. The time is in 24-hour format and is based on the server time zone, not the client time zone.
Recurrence	How often to create a deployment job based on this schedule: • One time—Run this job once on the day specified as the start date at the specified start time. • Hourly—Run this job on an hourly schedule. Specify the number of hours between deployment jobs. • Daily—Run this job on a daily schedule. Specify the number of days between deployment jobs. • Weekly—Run this job on the specified days of the week. • Monthly—Run this job on a monthly schedule. Select the day of the month to run the job, and the number of months between deployment jobs.
Run Indefinitely End Date and Time	The expiration date and time for the schedule. Deployment jobs are not created after this time. Select Run Indefinitely if you do not want the schedule to expire.

Table 8-15 Schedule Dialog Box (continued)

Element	Description
Devices To Deploy Group	
	This table lists the devices that are included in the deployment job. To add devices to the list, or to remove them from it, click **Add devices**, which opens the Add Other Devices dialog box (see Add Other Devices Dialog Box, page 8-58).
	If Security Manager is configured to use user-login credentials for accessing devices, your username and password are captured during schedule creation. If you change your password, you will need to recreate the schedule.

Add Other Devices Dialog Box

Use the Add Other Devices dialog box to select devices for the deployment job or schedule. The devices in the list might not have active policy changes. When you are creating a job, you might want to add devices that do not have policy changes if a device was manually modified and you want to return the device to its previous configuration (the configuration stored in the Security Manager database).

- Select the devices to include in the job or schedule in the Available Devices list and click **>>** to move the devices to the Selected Devices list.
- To remove devices, select them in the Selected Devices list and click **<<**.

Navigation Path

To open this dialog box, do one of the following:

- (Non-Workflow mode) From the Deploy Saved Changes dialog box, click **Add other devices**. See Deploying Configurations in Non-Workflow Mode, page 8-28.
- (Workflow mode) From the Deployment—Create or Edit a Job Dialog Box, click **Add other devices**. See Creating and Editing Deployment Jobs, page 8-35.
- (All modes) From the Schedule Dialog Box, click **Add devices**.

Related Topics

- Including Devices in Deployment Jobs or Schedules, page 8-8
- Creating or Editing Deployment Schedules, page 8-55
- Filtering Items in Selectors, page 1-45

Suspending or Resuming Deployment Schedules

You can suspend an active deployment schedule without discarding it and then reactivate it later when you want to resume creating jobs based on the schedule. This allows you to turn off a schedule temporarily.

Related Topics

- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Creating or Editing Deployment Schedules, page 8-55

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Schedules** tab if it is not active (see Deployment Schedules Tab, Deployment Manager, page 8-21).

Step 2 Do one of the following:

- To suspend an active schedule, select it and click **Suspend**.
- To resume a suspended schedule, select it and click **Resume**.

Adding Configuration Versions from a Device to the Configuration Archive

The Configuration Archive is updated with a new configuration version any time a configuration is deployed to the device or a file, including when you roll back a configuration to a device.

You can also retrieve a configuration directly from the device to add to the Configuration Archive. This is useful when changes have been made directly to device configurations, which are called out-of-band changes.

Note You cannot retrieve configurations from devices that are managed by AUS and that have been configured with dynamic IP addresses.

This procedure will help you retrieve a configuration from a device and add it to the archive.

Related Topics

- Viewing and Comparing Archived Configuration Versions, page 8-59

Step 1 Select **Manage > Configuration Archive** to open the Configuration Archive (see Configuration Archive Window, page 8-23).

Step 2 In the Device selector, select the device from which you want to retrieve the configuration. The archived configurations appear in the right pane.

Step 3 Click **Add from Device**. Security Manager logs into the device, retrieves the running configuration, and adds it to the archive.

Viewing and Comparing Archived Configuration Versions

Using the Configuration Archive, you can view the previous configurations for a device, compare versions of the configuration, and view the transcripts related to configuration deployment. To open the Configuration Archive window, select **Manage > Configuration Archive**.

To view the configuration versions for a device, select the device in the device selector. All archived versions are listed in the right pane. You can do the following:

- To view a configuration, select it and click **View**, which opens the Config Version Viewer dialog box with the configuration displayed in the left pane (for information about the dialog box, see Configuration Version Viewer, page 8-60).

If there is more than one type of configuration available for the selected version, you can choose which type to view using the **Config Type** field. A Full version is a complete configuration, whereas a Delta version is just the commands that were different between this version and the device's previous full configuration. Delta configurations might include negative commands.

- To compare configurations, select one and click **View**. In the Config Version Viewer window, select the configuration you want to compare in the **Compare with Version** field. The second version appears in the right pane with differences color-coded according to the caption below the display area.
- To view the transcript associated with the deployment of a configuration, do one of the following:
  - From the Configuration Archive window, double-click the icon in the Transcript column for the desired configuration.
  - When viewing a configuration in the left pane of the Config Version Viewer dialog box, click **Transcript View**.

  A transcript is the log file of transactions between Security Manager and a device captured during a deployment or rollback operation. It includes commands sent and received between server and device from the time of the deployment or rollback request, but it does not include communication that occurs during the initial discovery phase of deployment, when Security Manager obtains the current configuration from the device. If rollback is unsuccessful, there might be a partial transcript generated depending on which stage rollback or deployment failed. The transcript is displayed in the Transcript Viewer window (see Viewing Deployment Transcripts, page 8-62).

You can configure the number of configuration versions to archive on the Configuration Archive settings page (see Configuration Archive Page, page 11-6).

Related Topics

- Adding Configuration Versions from a Device to the Configuration Archive, page 8-59

Configuration Version Viewer

Use the Config Version Viewer window (when opened from the Configuration Archive) to view previous configurations for a device and to compare them to other archived configurations. You can compare any version to any other version in the archive for a selected device. The selected version appears in the left pane, and you can select another version for comparison from the list on the upper right of this window. For more information on viewing and comparing versions, see Viewing and Comparing Archived Configuration Versions, page 8-59.

Navigation Path

Select **Manage > Configuration Archive**, select a device whose configuration you want to view, select the configuration, and click **View**.

Related Topics

- Configuration Archive Window, page 8-23
- Viewing Deployment Transcripts, page 8-62
- Viewing and Comparing Archived Configuration Versions, page 8-59
- Adding Configuration Versions from a Device to the Configuration Archive, page 8-59

Field Reference

Table 8-16 Configuration Version Viewer Window (Configuration Archive)

Element	Description
Version ID	The configuration version to display in the left pane: • Previous—Display the version in the sequence before the one currently selected. • Next—Display the version in the sequence after the one currently selected. • Last—Display the last version in the list. • Specific Date and Time—Display the version created on that date and time.
Compare with version	The configuration version to compare to the version selected in the left pane, if you want to compare versions. The configuration is displayed in the right pane, with differences summarized and color coded as explained by the caption below the pane.
Config Type	The types of configurations that are available for viewing. The types differ based on the type of device. The types might indicate Full or Delta, which have the following meaning: • Full Configuration—The full configuration for the selected device as saved in the Configuration Archive. You can compare full configurations for a device. • Delta Configuration—The file that is generated by Security Manager during deployment and that represents policy changes between the configuration selected in the Version ID field and the most recently deployed version. **Note** Configuration versions resulting from out-of-band changes (for example, in the CLI) can be added to Configuration Archive using **Add from Device**, but no delta configuration file is generated.
First Difference button	Moves the cursor to the first difference noted between the configuration versions.
Previous Difference button	Moves the cursor to the previous difference noted between the configuration versions.
Current Difference button	Using the cursor, focuses on the currently selected difference in the window.
Next Difference button	Moves the cursor to the next difference noted between the configuration versions.
Last Difference button	Moves the cursor to the last difference noted between the configuration versions.
Transcript View button	Click this button to open the transcript viewer window, which displays the device communication transcript associated with this configuration.
Print button	Click this button to print the configuration.

Viewing Deployment Transcripts

Use the Transcript Viewer window to view the record of messages exchanged between Security Manager and a device. A transcript is the log file of transactions between Security Manager and a device captured during a deployment or rollback operation. It includes commands sent and received between server and device from the time of the deployment or rollback request, but it does not include communication that occurs during the initial discovery phase of deployment, when Security Manager obtains the current configuration from the device. For more information, see Viewing and Comparing Archived Configuration Versions, page 8-59.

Navigation Path

- Configuration Archive—Select **Manage > Configuration Archive** to open the Configuration Archive, select the device for which you want to view a transcript and double-click the **Transcript** icon in the row for the desired configuration version.

  You can also click the **Transcript View** button from the Configuration Version Viewer window when examining an archived configuration (see Configuration Version Viewer, page 8-60).
- Deployment Manager—Select **Manage > Deployments** to open the Deployment Manager, select the deployment job that includes the desired device deployment, select the Details tab in the lower pane, and double-click the **Transcript** icon in the row for the desired device.

Related Topics

- Configuration Archive Window, page 8-23
- Deployment Manager Window, page 8-16

Field Reference

Table 8-17 Transcript Viewer Window

Element	Description
Version ID	The configuration version for which you are viewing transcripts: • Previous—Display the transcripts for the version in the sequence before the one currently selected. • Next—Display the transcripts for the version in the sequence after the one currently selected. • Last—Display the transcripts for the last version in the list. • Specific Date and Time—Display the transcripts for the version created on that date and time.
Transcript Type	The type of transcript that you want to view. Some configuration versions have more than one transcript associated with them. Use this field to select which transcript to view.
Transcript Window	Displays the selected transcript. You can select text and copy it to the clipboard (Ctrl+C) for pasting in a text editor.
View button	Click this button to display the related configuration in the Config Version Viewer window (see Configuration Version Viewer, page 8-60).
Print button	Click this button to print the transcript.

Rolling Back Configurations

After you deploy a new configuration to a device, you can roll back the configuration to an older version if you find that the new configuration does not work correctly. However, it is usually a better idea to fix the configuration in Security Manager and deploy the fixed configuration, because rolling back a configuration creates a situation where the configuration defined in Security Manager is not the same one running on the device. Roll back configurations only in extreme circumstances.

The following topics will help you better understand and use configuration rollback:

- Understanding Configuration Rollback, page 8-63
- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70
- Performing Rollback When Deploying to a File, page 8-71

Understanding Configuration Rollback

If you deploy configurations to devices using the Device method, either to deploy the configuration directly to the device or to an intermediate server, you can roll back the configuration to an older version if you find that the new configuration does not work correctly. You cannot roll back to a configuration that was deployed to a file.

Caution It is usually a better idea to fix the configuration in Security Manager and deploy the fixed configuration, because rolling back a configuration creates a situation where the configuration defined in Security Manager is not the same one running on the device. After rollback, you should rediscover policies on the device to make the device configuration and its configuration in Security Manager consistent. Roll back configurations only in extreme circumstances.

You can roll back configurations using these tools:

- Deployment Manager—You can roll back a deployment to the last good configuration if that configuration was deployed to the device rather than to a file. To open the Deployment Manager, select **Manage > Deployments**.
- Configuration Archive—You can roll back deployment to any archived configuration that was deployed to the device or that originated from the device. To open the Configuration Archive, select **Manage > Configuration Archive**.

When you roll back a configuration, Security Manager does the following:

- On PIX Firewalls and ASA and FWSM devices, Security Manager uses the **replace config** option on the device's SSL interface to perform the equivalent of a reload (xlates are cleared, IPsec tunnels are torn down, and so on).
- For devices running IOS 12.3(7)T or later, Security Manager uses the **configure replace** command to replace the running configuration with the contents of a configuration file. Support for this command is dependent on the IOS version installed on the device:
  - On devices running IOS 12.3(7)T or later, Security Manager copies the configuration file to the startup configuration before executing the **configure replace** command. If the configure replace operation fails, Security Manager issues the **reload** command to reload the operating system using the contents of the startup configuration. The reload command restarts the system, which might result in a temporary network outage.

  - On routers running a version prior to 12.3(7)T, Security Manager copies the configuration file to the startup configuration and issues the **reload** command, which restarts the system. Security Manager uses the TFTP server and directory specified in the Configuration Archive settings page (see Configuration Archive Page, page 11-6) when using this method.
- The rolled-back configuration becomes another archived version in the Configuration Archive for that device.

Tip Configuration rollback does not include user account policies. When you roll back a configuration, the existing state of user accounts is not changed. This helps ensure that users can continue to log into the device.

Special considerations apply to the rollback of certain device types and configurations. See the following sections for more information:

- Understanding Rollback for Devices in Multiple Context Mode, page 8-64
- Understanding Rollback for Failover Devices, page 8-65
- Understanding Rollback for Catalyst 6500/7600 Devices, page 8-65
- Understanding Rollback for IPS and IOS IPS, page 8-66
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70

Understanding Rollback for Devices in Multiple Context Mode

If the configuration of the system execution space to which you are rolling back specifies connectivity options to security contexts (for example, vlan config) and there is a mismatch between the configuration selected for rollback and the current running configurations of the security contexts, Security Manager might not be able to connect to the security contexts. In such cases. we recommend that you roll back configurations for the security contexts before rolling back a configuration for the system execution space.

If you roll back a configuration for the system execution space of a device in multiple context mode to one that includes a different set of security contexts, after rollback the security contexts on the device might not match the security contexts managed by Security Manager that appear in the Device selector.

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Understanding Rollback for Failover Devices

If you roll back a configuration for a security context that contains a failover policy, Security Manager initially disables failover in the system execution space and both devices become active. After the rollback is completed, the devices should return to their failover configuration.

If a switchover occurs during rollback or connectivity between the active and standby units is lost, copy the bootstrap configuration to the standby unit after rollback completes. For more information, see Bootstrap Configuration for LAN Failover Dialog Box, page 50-26.

Note Security Manager can proceed with rollback action if and only if the following conditions are met:

- Both the Primary unit and Secondary unit must be in Active state.
- If configured on a link, the link must be up.
- If configured on LAN, the interface must be up.

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Understanding Rollback for Catalyst 6500/7600 Devices

If you roll back a configuration to a Catalyst 6500/7600 device that specifies connectivity options to service modules (for example, vlan config) and there is a mismatch between the configuration selected for rollback and the current running configuration, Security Manager might not be able to connect to the service modules. We recommend that you roll back configurations for the service modules before rolling back a configuration to the Catalyst 6500/7600 chassis.

Thus, the proper order for performing rollback on Catalyst 6500/7600 devices is:

1. Security contexts.
2. Service modules.
3. Chassis.

We recommend performing rediscovery after the rollback operation is complete.

If you are rolling back an FWSM deployment and the system is configured to retrieve security certificates when adding devices, you might need to retrieve the certificate after the rollback operation is complete. This can be done using either of the following methods:

- Retrieving the certificate on a per-device basis from Device Properties.
- Configuring Security Manager to automatically retrieve certificates after rollback. To do this, select **Tools > Security Manager Administration > Device Communication**, then select **Retrieve while adding devices** in the PIX/ASA/FWSM Device Authentication Certificates field (in SSL Certificate Parameters).

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70

- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Understanding Rollback for IPS and IOS IPS

Note From version 4.17, Cisco Security Manager does not support FWSM, IPS, and PIX devices. In addition, from this release, Cisco Security Manager will not provide any enhancements.

Special considerations apply to the rollback of IPS devices and IOS IPS devices. For IPS devices and IOS IPS devices, rollback could possibly include rolling back sensor updates or signature updates. The reason for this is that for IPS devices and IOS IPS devices, Security Manager supports not only the management of configuration but also the support of image management in the form of manual and automatic upgrades and signature updates. Keep in mind that when you do a rollback, you are rolling back the configuration, not the sensor updates or signature updates. These updates are downgraded only if the configuration cannot be rolled back without downgrading the updates.

Rollback is accomplished through Configuration Archive. For IPS devices and IOS IPS devices, only the current configuration is archived. The current configuration for one device version (say, Version X) may not be valid for a different device version (say, Version Y). Security Manager rolls back a configuration of Version X to a sensor with Version Y as long as the configuration for X is valid for Y.

If the configuration for X is valid for Y, rollback proceeds and Security Manager displays a confirmation dialog box to you. If the configuration for X is not valid for Y, Security Manager displays a warning dialog box to you and provides you with the option of downgrading the sensor during rollback if such a downgrade will help accomplish the rollback.

Caution Downgrading an IPS device removes certain capabilities of the IPS device. For example, downgrading the engine prevents you from applying the latest signature updates. Operation of an IPS device without the latest signature updates diminishes the effectiveness of the IPS device.

For rollback of a deployment job, the warning dialog box contains one or more of the following types of warnings:

- Security Manager warns you about IPS devices that need to have their sensor version downgraded before a rollback can be performed.
- Security Manager warns you about IOS IPS devices whose signature level has changed. For these devices, only the non-IPS sections of the configuration can be rolled back.
- Security Manager warns you about IPS devices that must be downgraded more than one level, and as a result, Security Manager cannot do it. You must use the Cisco IPS CLI for such downgrades. The warning dialog box displays the version to which the device must be reimaged or downgraded.

Note The option of downgrading an IOS IPS device during rollback is not available, because IOS IPS devices do not support downgrade.

If the option of downgrading the sensor during rollback will not help accomplish the rollback, you receive an error message stating that rollback cannot occur and that you need to manually reinstall the image on the device to roll back. Only the update package most recently installed on a device can be downgraded, so downgrade does not help in the following cases:

- Rollback of a deployment (signature update) that involves downloading more than one update package to the device.
- Selection of an old deployment or configuration for rollback subsequent to which several upgrades occurred.
- Rollback of an upgrade that cannot be downgraded. Major, minor, and most service pack upgrades cannot be downgraded, as shown in Table 8-18 on page 8-67

For rollback of a configuration that requires a downgrade to a version prior to Cisco IPS 5.1(4), Security Manager does not support automatic downgrade. You must manually downgrade the device to the specified version and then proceed with rollback.

Table 8-18 Downgrade Support for Possible Sensor Upgrade Types

Upgrade Type	Downgrade Support
Major Upgrade	Downgrade is not supported.
Minor Upgrade	Downgrade is not supported.
Service Pack Update	Downgrade from Cisco IPS 5.1(4) onward is not supported.
Patch update	Downgrade is supported.
Signature Update	Downgrade is supported.
Engine Update	Downgrade is supported.
Repackage (applicable to major, minor, and service pack updates).	Repackages for service packs prior to 5.1(4) can be downgraded.

Caution Outbreak Prevention updates on a particular device may be lost if that device is downgraded.

During rollback, if Security Manager discovers that there have been out-of-band changes to the device that prevent rollback, you will receive an error message stating that rollback is prevented.

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70

Commands that Can Cause Conflicts after Rollback

The following commands can potentially cause conflicts after rollback is performed:

- **http server enable** *port*
  http *ip_address net_mask interface_name*

  Applicable only to security contexts (not the system execution space).
- **allocate-interface** {*physical_interface* | *subinterface*} **[map_name] [visible | invisible]**

  Applicable only to the system execution space under the context subcommand.
- **config-url** *diskX:/path/filename*

  Applicable only to the system execution space under the context subcommand.
- **join -failover-group** *group_number*

Applicable only for active/active failover and only to the system execution space under the context subcommand. The failover group defaults to group 1 if not specified.

- **failover**

  Applicable only to the system execution space. Enabling failover causes configuration synchronization to trigger between peers.

- **failover lan enable**

  Applicable only to the system execution space. If this command is omitted, this implies serial cable failover on a PIX platform or warrants an incomplete failover configuration warning on ASA and FWSM.

- **failover lan unit** {*primary* | *secondary*}

  Applicable only to the system execution space. If this command is not specified, both units are secondary by default. If rollback takes place on the wrong unit, both can become primary, which impacts which unit becomes active initially.

- **failover group** *group_number*

  Applicable only to the system execution space. This command enables active/active failover. If this command is omitted, active/standby is enabled.

- **preempt** *delay*

  Applicable only to the system execution space and under the failover group subcommand to force which failover group becomes active if both units are booted up at the same time, or the primary does not boot up within the delay specified.

- **monitor-interface** *interface_name*

  Applicable only to security contexts and used to enable health monitoring of critical interfaces. If this interface is 'bounced' or fails, a switchover could occur.

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Commands to Recover from Failover Misconfiguration after Rollback

If a switchover happens during rollback and the two units are no longer synchronized, you might need to use the following commands to recover:

- **failover active** *group_number*
- **failover reset** *group_number*
- **failover reload-standby**
- **clear configure failover**

For more information on these commands, please refer to the command reference for your security appliance.

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Using Rollback to Deploy Archived Configurations, page 8-70

- Commands that Can Cause Conflicts after Rollback, page 8-67

Rolling Back Configurations to Devices Using the Deployment Manager

If you deploy configurations to devices and then determine that there is something wrong with the new configurations, you can revert to and deploy the previous configurations for those devices. You cannot roll back to a previous configuration if there are no previous configurations in the Configuration Archive.

You can roll back configurations only to configurations that were deployed to the device, not to a file. For information on how to roll back a configuration that was deployed to a file, see Performing Rollback When Deploying to a File, page 8-71.

You can also use the Configuration Archive tool to roll back to any configuration archived from a device. For more information, see Using Rollback to Deploy Archived Configurations, page 8-70.

Caution Roll back configurations only in extreme circumstances. It is usually a better idea to fix the configuration in Security Manager and deploy the fixed configuration, because rolling back a configuration creates a situation where the configuration defined in Security Manager is not the same one running on the device. After rollback, you should rediscover policies on the device to make the device configuration and its configuration in Security Manager consistent. Roll back configurations only in extreme circumstances. Before proceeding, read the following topics.

- Understanding Configuration Rollback, page 8-63
- Understanding Rollback for Devices in Multiple Context Mode, page 8-64
- Understanding Rollback for Failover Devices, page 8-65
- Understanding Rollback for Catalyst 6500/7600 Devices, page 8-65
- Understanding Rollback for IPS and IOS IPS, page 8-66
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Before You Begin

When you roll back a configuration, the action is not done as part of an activity or configuration session, which means the device is not locked. Thus, it is possible that two users might roll back configurations simultaneously on a device, which can generate unexpected problems. Before rolling back a configuration, ensure that there are no active deployment jobs for the device listed in the Deployment Manager window.

Related Topics

- Viewing Deployment Status and History for Jobs and Schedules, page 8-26
- Job States in Non-Workflow Mode, page 8-4
- Job States in Workflow Mode, page 8-6

Step 1 Click the **Deployment Manager** button in the Main toolbar. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the deployment job (which must be in the Deployed or Failed states) and click **Rollback**.

The Rollback a Job dialog box opens. The dialog box lists all of the devices included in the job, including the name of the device, the deployment method (file or device), the status of the previous deployment, and the name of the deployment job that last updated the device.

Step 3 Select the devices for which you want to roll back configurations by checking the check box in the Selection column. You can select only devices that used the deploy to device method. By default, all the devices with the status Succeeded are selected.

You can view the configuration that will be deployed to a device by highlighting the row for a device and clicking the **Preview Config** button. You can compare it to the last deployed configuration or the current running configuration. For more information, see Previewing Configurations, page 8-44.

Step 4 Click **OK**. You are asked to confirm the action.

Step 5 (Optional) To make the configuration defined in Security Manager consistent with the one running on the device, rediscover the device policies as described in Discovering Policies on Devices Already in Security Manager, page 5-15.

Using Rollback to Deploy Archived Configurations

You can roll back any configuration version from Configuration Archive to the device for which it is archived, provided that the configuration was deployed to the device or originated from the device. The rolled-back configuration then becomes another archived version in the list for that device. For information on how to roll back a configuration that was deployed to a file, see Performing Rollback When Deploying to a File, page 8-71.

Before You Begin

Tip When you roll back a configuration, the action is not done as part of an activity or configuration session, which means the device is not locked. Thus, it is possible that two users might roll back configurations simultaneously on a device, which can generate unexpected problems. Before rolling back a configuration, check the Deployment Manager to ensure that there are no active deployment jobs for the device (select **Manage > Deployments**).

Roll back configurations only in extreme circumstances. Before rolling back configurations, carefully read these topics:

- Understanding Configuration Rollback, page 8-63
- Understanding Rollback for Devices in Multiple Context Mode, page 8-64
- Understanding Rollback for Failover Devices, page 8-65
- Understanding Rollback for Catalyst 6500/7600 Devices, page 8-65
- Understanding Rollback for IPS and IOS IPS, page 8-66
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Related Topics

- Rolling Back Configurations to Devices Using the Deployment Manager, page 8-69
- Adding Configuration Versions from a Device to the Configuration Archive, page 8-59

- Chapter 8, "Managing Deployment"
- Viewing and Comparing Archived Configuration Versions, page 8-59

Step 1 Select **Manage > Configuration Archive** to open the Configuration Archive (see Configuration Archive Window, page 8-23).

Step 2 In the Device selector, select the device for which you want to roll back to a different configuration version. The archived configurations appear in the right pane.

Step 3 Select the configuration version to which you want to roll back. You can roll back only to a configuration that was deployed to the device or that originated from the device. You cannot roll back to a configuration that was deployed to a file.

Tip To view the configuration version before rollback, click **View**.

Step 4 Click **Rollback** to deploy the selected configuration version to the device. A progress box appears, followed by a notification message when the configuration version is successfully deployed.

Step 5 (Optional) To make the configuration defined in Security Manager consistent with the one running on the device, rediscover the device policies as described in Discovering Policies on Devices Already in Security Manager, page 5-15.

However, it is usually better to correct the policies for the device and to then redeploy the updated configuration. This preserves your changes and shared-policy configuration for the device, which would otherwise be removed if you rediscover policies.

Performing Rollback When Deploying to a File

You cannot directly perform rollback when deploying to a file instead of to a device. Use this procedure to revert to a previously stored configuration when deploying to file.

Related Topics

- Understanding Configuration Rollback, page 8-63
- Understanding Rollback for Devices in Multiple Context Mode, page 8-64
- Understanding Rollback for Failover Devices, page 8-65
- Understanding Rollback for Catalyst 6500/7600 Devices, page 8-65
- Understanding Rollback for IPS and IOS IPS, page 8-66
- Commands that Can Cause Conflicts after Rollback, page 8-67
- Commands to Recover from Failover Misconfiguration after Rollback, page 8-68

Step 1 Select **Manage > Configuration Archive** to open the Configuration Archive (see Configuration Archive Window, page 8-23).

Step 2 In the Device selector, select the device for which you want to roll back to a different configuration version. The archived configurations appear in the right pane.

Step 3 Select the configuration version to which you want to roll back and click **View**.

Step 4 In the Configuration Version Viewer window, make sure the Config Type is set to Full.

Step 5 Click in the left-hand pane, then press Ctrl+A followed by Ctrl+C to copy the selected configuration to the Windows clipboard.

Step 6 Open a text editor such as NotePad, then press Ctrl+V to paste the contents of the clipboard into the text file.

Step 7 Save the file. You can use this file to perform manual rollback.

Troubleshooting Device Communication and Deployment

One of the more likely areas where you can run into problems is with actions where Security Manager must log into a device. These types of actions include policy discovery and deployment using live devices, or actions that involve retrieving information from a device.

The key point to remember is that the communication pathway is from the Security Manager server to the device; the workstation on which you are running the Security Manager client is not involved in device communication (unless the server is installed on the same machine, of course). The Security Manager server must have a network pathway to the device as well as the correct credentials and certificates to authenticate with the device for communication to be successful.

The following topics can help you troubleshoot general device communication and policy deployment problems:

- Testing Device Connectivity, page 9-1
- Managing Device Communication Settings and Certificates, page 9-4
- Resolving Red X Marks in the Device Selector, page 9-9
- Troubleshooting Deployment, page 9-9

Testing Device Connectivity

Security Manager must be able to connect to and log into a device in order to manage it. You can test whether Security Manager can use the credentials and transport method you have defined within Security Manager for this purpose.

You can test connectivity only for devices that have static IP addresses. You also cannot test connectivity for devices that use Token Management Server (TMS) as the transport protocol.

If you add a device from the network or from an inventory file to the inventory, Security Manager tests connectivity automatically.

You can manually test device connectivity for any device in the inventory or for new devices that you are adding manually. The following procedure describes how to test connectivity for devices that are already in the inventory. When adding devices manually, click **Test Connectivity** on the Device Credentials page of the New Device wizard to perform the test described below. For more information on adding devices manually, see Adding Devices by Manual Definition, page 3-26.

Before You Begin

Security Manager uses the settings on the Device Communication page to determine the connection timeout, how often to retry the connection, the transport protocol, and which credentials to use. To configure these settings, select **Tools > Security Manager Administration** and select **Device Communication** from the table of contents.

Related Topics

- Understanding the Device View, page 3-1
- Viewing or Changing Device Properties, page 3-41
- Device Communication Page, page 11-21

Step 1 In Device view, do one of the following in the Device selector to open the Device Properties dialog box:

- Double-click a device.
- Right-click a device and select **Device Properties**.
- Select a device and select **Tools > Device Properties**.

Step 2 Select **Credentials** from the table of contents.

Step 3 Click **Test Connectivity**.

The Device Connectivity Test dialog box opens and displays the progress of the test, including the protocol being used (see Device Connectivity Test Dialog Box, page 9-3). You can abort the test while it is running. When the test is finished, click **Details** to see:

- For successful tests, the output of the **show version** command or the **getVersion** command (for IPS Sensors and Cisco IOS IPS Sensors). You can select the text, press Ctrl+C to copy the text to the clipboard, and then paste it into another file for later analysis.
- For unsuccessful tests, the error information. Some common problems are:
  - The username or password is incorrect.
  - The wrong protocol is selected. For example, the device might not be configured to respond to the selected protocol.
  - The device is not configured to accept connections correctly. Ensure that at least one supported protocol is configured.
  - The wrong operating system is specified for the device (for example, you specified PIX for an ASA device).
  - If you are using ACS authentication and the connection to the device is completed, you can get errors when Security Manager tries to obtain version information if you do not have Control authorization.
  - There might be general network configuration problems. Test connectivity to the device from outside of Security Manager. Look for hardware, media, and booting errors, excessive traffic causing queues to overflow, duplicate MAC or IP addresses on the device, physical discrepancies, such as link, duplex, and speed mismatch, or logical discrepancies, such as VLAN and VTP inconsistencies or ATM network misconfiguration.

Device Connectivity Test Dialog Box

Use the Device Connectivity Test dialog box to view whether Security Manager can contact the device using the configured credentials.

Navigation Path

To start the device connectivity test, click **Test Connectivity** from the Credentials page in one of these areas:

- New Device wizard when adding a device manually. See Adding Devices by Manual Definition, page 3-26.
- Device Properties. To open the page, double-click a device in the Device selector or select **Tools > Device Properties**.

The connectivity test is done automatically when you click **Next** or **Finish** on the Credentials page when adding a device from the network.

Related Topics

- Testing Device Connectivity, page 9-1
- Device Credentials Page, page 3-45
- Viewing or Changing Device Properties, page 3-41

Field Reference

Table 9-1 Device Connectivity Test Dialog Box

Element	Description
Connectivity Protocol	The transport protocol being used to log into the device. Security Manager uses the protocol specified in the device properties for the device, which is usually the default protocol configured on the Device Communications page (see Device Communication Page, page 11-21).
Connectivity Status	Displays the status of the test and the time elapsed since the start of the test.
Details button	Click this button to display detailed information about the result of the test. • Passed tests—The details display the output of the **show version** command for PIX Firewall, Adaptive Security Appliances (ASA), Firewall Service Modules (FWSM), Cisco IOS routers, and VPN Services Modules (VPNSM), or the output of the **getVersion** command for IPS Sensors and Cisco IOS IPS Sensors. You can copy the command output and paste it into a file for analysis. • Failed tests—The detailed error message.
Abort button	Stops the connectivity test before it is completed.

Managing Device Communication Settings and Certificates

If you discover device inventory and policies directly from devices, or deploy configurations to devices rather than to files, you must configure Security Manager to use the transport protocols that your devices use. For some device types, only one transport protocol is supported, so you do not need to make a choice. For other devices, such as Cisco IOS routers, you have options concerning the protocols you use.

Security Manager has default settings for transport protocols that are the most-used protocols for each device type. To change these settings, select **Tools > Security Manager Administration** and select **Device Communication** from the table of contents (see Device Communication Page, page 11-21).

For most users, the communication settings that require management are the certificates used for SSL (HTTPS) connections and the public keys used for SSH connections. You might update the certificates and keys on the device, which would leave Security Manager holding an outdated copy.

The following topics provide more information about managing certificates and keys, and how to troubleshooting device communications:

- SSL certificates—You can configure Security Manager to automatically replace certificates using the ones obtained from the device on the Device Communication page. If you decide to manually manage the SSL certificate store, see Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5. The following topics provide more information about certificate errors:
  - Security Certificate Rejected When Discovering Device, page 9-6
  - Invalid Certificate Error During Device Discovery, page 9-7
  - Managing IPS Certificates, page 44-10

Tip Ensure that all PIX Firewalls and Adaptive Security Appliances that you intend to manage with Security Manager have a 3DES/AES license. See Understanding Device Communication Requirements, page 2-1.

- SSH Public Keys—By default, Security Manager replaces public keys with the new ones obtained during SSH connections. If you have problems with SSH communications, see Troubleshooting SSH Connection Problems, page 9-7.
- General Device Communication Troubleshooting—For other problems you might encounter, see Troubleshooting Device Communication Failures, page 9-8.

Multiple Certificate Authentication Support

Beginning from version 4.13, the Cisco Security Manager supports ASA 9.7.1 feature of multiple certificate authentication for VPN connectivity. ASA, in its release 9.7.1, has introduced multiple certificate authentication support to its VPN client customers. Thus, the client can authenticate remote VPN users with two client certificates. The two client certificate could be a combination of one user certificate and one machine certificate, or two user certificates. Two machine certificate authentication is not supported for security considerations. The multiple certificate authentication works for both SSL VPN and IPsec VPN.

In Cisco Security Manager 4.13, to enable the multiple certificate authentication support, you are required to appropriately specify the AAA authentication method (see AAA Tab (Connection Profiles), page 31-13) and configure the DAP policy (see Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication, page 32-38).

Manually Adding SSL Certificates for Devices that Use HTTPS Communications

Note In addition to the techniques described in this topic, for IPS devices you can use the IPS Certificates utility to manage the certificates in Security Manager's certificate data store. For more information, see Managing IPS Certificates, page 44-10.

When you use SSL (HTTPS) as the transport protocol for communicating with IPS, PIX, ASA, or FWSM devices, or Cisco IOS routers, you can configure Security Manager to automatically retrieve the device authentication certificate when adding the device (see Device Communication Page, page 11-21).

Tip Having an accurate certificate is required for successful HTTPS communications; Security Manager cannot communicate with the device without the correct certificate, which prevents configuration deployment. When using self-signed certificates, the device might create a new certificate if Security Manager attempts to access it using the wrong certificate. Thus, it is best to configure Security Manager to always retrieve the certificate from the device.

Instead of having Security Manager automatically retrieve the certificates, you can manually add them to increase the level of network security. On the Device Communication page, you would configure the device authentication setting for the device type as **Manually add certificates**.

The easiest way to manually update the certificate for a device is to retrieve it from the device. Right-click the device and select **Device Properties**. Click **Credentials** to open the Credentials page, and then click **Retrieve From Device** to the right of the **Authentication Certificate Thumbprint** field. Security Manager retrieves the certificate and prompts you to accept it. You might need to do this if you encounter certificate problems during configuration deployment. (You can also type or paste the certificate into this field.)

You can also manually type in, or copy and paste, the certificate thumbprint without having Security Manager log into the device. Use the following procedure to manually enter the SSL certificate thumbprint for a device if you configured that device type to require manually added certificates.

Tip Security Manager allows you to generate a 2048-bit self-signed certificate under Megamenu > Server Administration > Server > Security > Single Server Management > Certificate Setup.

Tip To reach the Megamenu, double-click the Cisco Security Manager icon on your server desktop and log on. Another way to reach the Megamenu is as follows: Windows > Start > All Programs > Cisco Security Manager > Cisco Security Manager > [log on]. This second navigation path may differ slightly, depending upon how you have personalized your Windows Server installation.

Before You Begin

Obtain the certificate thumbprint (a hexadecimal string) for the device.

Tip If the thumbprint is not readily available, you can copy it from the error message that is displayed when you add the device from the network or from an export file.

Step 1 Select **Tools > Security Manager Administration** and select **Device Communication** from the table of contents to open the Device Communication page (see Device Communication Page, page 11-21).

Step 2 Click **Add Certificate** to open the Add Certificate dialog box (see Add Certificate Dialog Box, page 11-24).

Step 3 Enter the DNS hostname or IP Address of the device, the certificate thumbprint in hexadecimal format, and click **OK**. The thumbprint is added to the certificate store.

Tip To erase an existing thumbprint, leave the Certificate Thumbprint field empty.

Security Certificate Rejected When Discovering Device

If an error occurs when you attempt to discover a device, and the error message states that the security certificate received from the device was rejected, you need to update the certificate. You can do this using one of the following methods:

- For IPS devices only, select **Manage > IPS > IPS Certificates** and synchronize the certificates. You might also need to regenerate the certificate. For more information, see Managing IPS Certificates, page 44-10.
- Manually enter the thumbprint required by the certificate by doing one of the following:
  - Select **Tools > Security Manager Administration > Device Communication**. Click **Add Certificate**, enter the IP address of the device, then copy and paste the thumbprint displayed in the error message into the Certificate Thumbprint field.
  - Right-click the device and select **Device Properties > Credentials**. Copy and paste the thumbprint displayed in the error message into the Authentication Certificate Thumbprint field.

  You must manually enter the thumbprint whenever you add a new device using the Add New Device or Add From Configuration File options and when you perform rediscovery. It is not required when you add a new device using the Add New Device From Network or Add Device From File options.
- Configure the SSL certificate settings to automatically retrieve the certificate when adding devices. You can select different settings for IPS, router, and ASA/PIX/FWSM devices. To configure these settings, select **Tools > Security Manager Administration > Device Communication**, and look at the **SSL Certificate Parameters** group.

Related Topics

- Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5
- Adding Devices to the Device Inventory, page 3-7
- Chapter 2, “Preparing Devices for Management”
- Device Communication Page, page 11-21
- Device Credentials Page, page 3-45

Invalid Certificate Error During Device Discovery

If the time settings on the device and Security Manager are not in synchronization, when you try to discover policies on a device (adding it to the inventory or rediscovering policies on a device already in the inventory), an error message might state that the certificate is not yet valid.

When the time set on the Security Manager server is lagging behind the time set on the device, Security Manager cannot validate the device certificate if the start time of the validity period is ahead of the Security Manager time setting. Even if the time zones configured on the device and Security Manager are the same, the invalid certificate error occurs if the daylight saving time (summertime) settings are different. To resolve this problem, make sure that the daylight saving time settings are the same on the device and Security Manager, regardless of whether the time zone is the same. After setting the daylight saving time, synchronize the clock on the device with Security Manager so that both of them display the same time.

To obtain best results, we recommend that you set the same time zone on the device and Security Manager, and modify the time zone after you discover the certificates at a later time, if necessary.

Related Topics

- Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5
- Managing IPS Certificates, page 44-10
- Adding Devices to the Device Inventory, page 3-7
- Chapter 2, “Preparing Devices for Management”

Troubleshooting SSH Connection Problems

For devices that use SSH as the transport protocol, Security Manager automatically detects the appropriate SSH version (1.5 or 2) to use with each device. During SSH version 2 connections, Security Manager automatically negotiates encryption algorithms or ciphers with the device. Security Manager also automatically overwrites the SSH public key for the device if the key changes. Thus, you typically will not run into SSH connection problems.

If you do have SSH connection problems, consider these fixes:

- If the public key on the device changed, and SSH connections are failing due to a key problem, remove the key for the device from the Program Files/CSCOpx/MDC/be/tmp/.shh/known_hosts file on the Security Manager server and retry the operation.
- Security Manager uses 3DES (Data Encryption Standard) as the default encryption algorithm. If this is not the correct algorithm for your devices, either change the configuration of your devices, or update the Program Files/MDC/athena/config/DCS.properties file to indicate the correct algorithm on the DCS.ssh.encipher property. (Contact Cisco TAC if you need more help). You must restart the Security Manager daemon manager if you change this file.

Related Topics

- Chapter 2, “Preparing Devices for Management”
- Device Communication Page, page 11-21
- Device Credentials Page, page 3-45

Troubleshooting Device Communication Failures

If Security Manager fails to communicate with a device, for example, by failing to log into it, during discovery, deployment, or other actions, look at these areas to identify and resolve the problem:

- Ensure the device is operational.
- Check which transport protocol is selected. You must select a protocol that the device is configured to accept. For most devices, the protocol is selected on the Device Properties General page (select **Tools > Device Properties > General**). For IPS devices, the IPS RDEP mode is selected on the device properties Credentials page.

  For IOS devices that do not have a K8 or K9 crypto image, ensure that you select Telnet as the protocol.

  Some methods of adding devices also allow you to select a non-default transport protocol. To configure the default transport protocols for classes of devices, select **Tools > Security Manager Administration > Device Communications**.
- On the Device Properties General page, ensure that the hostname, domain name, and IP address are correct. Keep in mind that the Hostname and Accounts and Credentials policies for the device define the actual names and credentials that get configured on the device. However, the policies are not used for device communication. If you make changes to the policies that affect the credentials you are using for device communication, you must also manually update the device properties.
- Make sure DNS names can be resolved from the Security Manager server. You might need to fix the DNS settings on the server.
- Check the credentials for the device in Security Manager and ensure that they are correct and that there is a route between the server and device. Right-click the device, select **Device Properties**, select the Credentials tab, and click the **Test Connectivity** button. If the connection fails, check error messages to determine whether the problem is connectivity or credentials. Update the credentials in the device properties if necessary.

  When adding new devices the credentials are defined within the New Device wizard if your method of adding the device requires credentials. Keep the following in mind:

  - The primary credentials are used for SSH and Telnet connections.
  - The HTTP/HTTPS credentials are used for HTTP and SSL connections unless you select **Use Primary Credentials**, in which case the primary credentials are also used for these connections.
- Beginning with version 4.11, Security Manager does not support the device SSL Certificates using MD5 algorithms. If the device SSL uses MD5 algorithms, Security Manager throws up an error when you try to add the device to Security Manager. This happens because JRE, by default, disables the MD5 algorithms due to security vulnerability. To resolve this, you must use higher encryption algorithm for device SSL certificate.

Related Topics

- Adding Devices to the Device Inventory, page 3-7
- Understanding Device Communication Requirements, page 2-1
- Chapter 2, “Preparing Devices for Management”
- Device Communication Page, page 11-21
- Device Credentials Page, page 3-45

Resolving Red X Marks in the Device Selector

If a device is marked with a red X in the device selector in Device view, it means that the Auto Update Server (AUS) or Configuration Engine server assignment for the device was lost during an upgrade from a Security Manager release prior to 3.2.0. AUS and Configuration Engines are not migrated during an upgrade from 3.1.x, and devices managed by them need to be reassigned to them after the upgrade using the following procedure.

Step 1 Do one of the following in Device view:

- From the Device selector, right-click a device with a red X icon, then select **Update Server Info**.
- Click any red X icon in the device selection tree. A warning message is displayed stating that AUS and Configuration Engine information was not migrated after the upgrade process. Click **Yes** to add these servers manually.

The Device Server Assignment dialog box opens.

Step 2 From the Available Devices list, select all the devices that use the same AUS or Configuration Engine server and click **>>** to move them to the selected list. The Available Devices list includes all devices that are managed by AUS or Configuration Engine that are marked with a red X.

Step 3 Select the AUS or Configuration Engine that manages the selected devices from the Server list. If the correct server is not listed, select **+ Add Server...** to add it to the inventory using the Server Properties Dialog Box, page 3-38.

For more information on adding AUS or Configuration Engine servers to the inventory, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36.

Step 4 Repeat the process until no device is marked with a red X.

Troubleshooting Deployment

The deployment process is one of the most likely areas in which you will encounter problems when using Security Manager. There are many different processes involved in deployment that influence whether the deployment job is successful:

- Security Manager itself.
- The stability and availability of your network, including links to remotely-managed devices.
- Any bugs inherent in the versions of the operating systems that you are using on your network devices that affect the commands Security Manager is trying to deploy (Security Manager is not immune to these bugs).
- The licenses you have enabled on the device, because many security commands require specific device licenses.
- The specific features supported by your devices, which Security Manager cannot always determine ahead of time. For example, some platforms support features only if the device has a certain minimum RAM, and some interface settings are available for specific interface cards only.
- The correct functioning of interim applications such as AUS, Configuration Engine, or your TMS server.

If you encounter deployment failures, examine the messages in the deployment status window carefully. In addition, the following topics address some of the problems you might encounter:

- Changing How Security Manager Responds to Device Messages, page 9-10
- Memory Violation Deployment Errors for ASA 8.3+ Devices, page 9-12
- Security Manager Unable to Communicate With Device After Deployment, page 9-12
- Updating VPNs That Include Routing Processes, page 9-13
- Mixing Deployment Methods with Router and VPN Policies, page 9-14
- Deployment Failures for Routers, page 9-14
- Deployment Failures for Catalyst Switches and Service Modules, page 9-16
- Deployment Failures to Devices Managed by AUS, page 9-18
- Troubleshooting the Setup of Configuration Engine-Managed Devices, page 9-18

Changing How Security Manager Responds to Device Messages

Security Manager has built-in responses to many of the response messages that can be encountered when configuring a device. You might find that messages Security Manager treats as errors are messages that you want to ignore or treat as informational. Although you can configure your deployment jobs to ignore errors, you might instead want to update Security Manager to treat specific messages differently using a properties file.

It is important to understand that setting the properties file to ignore the error is not always sufficient. Deployment can fail because the **Allow Download on Error** check box (located on the **Tools > Security Manager Administration > Deployment** page) is deselected by default. The following table provides details about how Security Manager behaves when an error occurs during deployment, the **Allow Download on Error** option is either selected or deselected, and the **Save Changes Permanently on Device** option is selected or deselected.

Table 9-2 Deployment Device Error Handling for SSL and SSH on PIX Firewall, ASA, and Cisco IOS Routers

Allow Download on Error	Error Occurred	Error Ignored Using Warning Expression	Deployment Status	Write Memory Done
Selected	Yes	No	Failed	Based on whether **Save Changes Permanently on Device** is selected.
Selected	Yes	Yes	Success	Based on whether **Save Changes Permanently on Device** is selected.
Selected	No	Not applicable	Success	Based on whether **Save Changes Permanently on Device** is selected.
Deselected	Yes	No	Failed (Deploy not Completed message)	No.

Table 9-2 Deployment Device Error Handling for SSL and SSH on PIX Firewall, ASA, and Cisco IOS Routers (continued)

Allow Download on Error	Error Occurred	Error Ignored Using Warning Expression	Deployment Status	Write Memory Done
Deselected	Yes	Yes	SSL (ASA, PIX, IOS devices)—Failed SSH (IOS devices—Success	SSL—No. SSH (IOS devices)—Based on whether **Save Changes Permanently on Device** is selected.
Deselected	No	Not applicable	Success	Based on whether **Save Changes Permanently on Device** is selected.

Note On Cisco IOS routers using the SSL protocol, deployment on devices stops on command syntax errors. It does not stop when configuration-related errors occur.

To change how Security Manager treats a message, you need to update the DCS.properties file in \CSCOpx\MDC\athena\config folder in the installation directory (usually c:\Program Files). Use a text editor such as NotePad to update the file.

It is easiest to determine the message you want to ignore by looking at the transcript of a deployment job that encountered the error using the following procedure.

Related Topics

- Viewing Deployment Status and History for Jobs and Schedules, page 8-26

Step 1 Click the **Deployment Manager** button in the Main toolbar.

The Deployment Manager window appears. Click the **Deployment Jobs** tab if it is not active.

Step 2 Select the job with the error message.

Step 3 Click the **Transcript** button in the Deployment Details tab to open the transcript.

Step 4 Identify the error text that you want to ignore.

Step 5 Locate the appropriate warning expressions property in the DCS.properties file. For example, for PIX devices the property is called **dev.pix.warningExpressions**, whereas for IOS devices the property is called **dev.ios.warningExpressions**.

Tip Conversely, you can make device responses that are not tagged with the Error prefix to appear as error messages. To do this, add the message to the Error Expressions list (for example, **dev.pix.ErrorExpressions**).

Step 6 Add the error text to the warning expressions list. The warning message should be a generic regular expression string. Except for the last expression, you must delimit all expressions with "$\". For example, if the message you want to ignore is "Enter a public key as a hexadecimal number," enter the following string:

.*Enter a public key as a hexadecimal number .*$

Step 7 Restart the CiscoWorks Daemon Manager.

Memory Violation Deployment Errors for ASA 8.3+ Devices

ASA Software release 8.3+ requires significantly more device memory than previous versions of the ASA software. If you upgrade an ASA device that does not meet the minimum memory requirements, the upgrade process notifies you of the problem and the device regularly sends syslog messages until the minimum memory requirement is met.

Because an ASA device that does not meet the minimum memory requirements can function poorly, Security Manager does not deploy configurations to these devices, although you are allowed to add the device to the inventory and discover policies from it. However, if you try to deploy policies to the device before you add memory, you get a deployment error stating that the device does not meet the minimum memory requirements and deployment fails.

The best way to resolve the error is to add memory to the device. For information about ASA devices and memory upgrade possibilities, see *http://www.cisco.com/go/asa*.

Alternatively, you can downgrade the ASA software version, in which case you should delete the device from the inventory, then add it back to the inventory and discover policies.

Error While Attempting to Remove Unreferenced Object

If you enable the Remove Unreferenced Object Groups from Device option on the **Tools > Security Manager Administration > Deployment** page, Security Manager will remove objects during deployment that are not used in any policies managed or discovered by Security Manager. If any policy that is NOT discovered or managed by Security Manager is using such an object, Security Manager will still attempt to delete that object during deployment. In such cases, deployment will fail with a transcript error indicating that it was unable to delete the object as the object is being used. To successfully deploy, disable the Remove Unreferenced Object Groups from Device option.

Security Manager Unable to Communicate With Device After Deployment

There are a number of policies that you can configure in Security Manager that prevent access to a device. That is the point of security, ensuring that unwanted hosts cannot enter your network or network devices.

However, you can inadvertently lock the Security Manager server out of a device, making it impossible for Security Manager to deploy configurations to it or manage it. If you find that after a deployment, Security Manager can no longer contact a device, and you have already checked that the device is up and running and otherwise functioning normally, look into the following policies to see if they are causing the lock-out:

- **Firewall > Access Rules, or Firewall > Zone Based Firewall Rules** —If you use these policies, the rules must allow management traffic from the Security Manager server. Consider allowing at least HTTP, HTTPS, SSH, and Telnet. Consider creating a shared policy that defines the required access for Security Manager and applying it to all devices. Keep in mind that if you create any rules in these policies, an implicit rule is added to the end of the policy that denies any traffic that is not explicitly allowed.

- **NAT policies**—Make sure that you are not using a local address on the device as the original address to be translated. Translating this address might result in translating the management traffic sent between Security Manager and the device, causing the interruption.
- **Device Access policies on routers**—Security Manager might loose contact with a device after you unassign a device access policy from the device and redeploy it. Device access policies can be used to define the enable password for accessing the device. If you unassign this policy and redeploy, the password is removed from the device. In such cases, the device typically reverts to the default password. However, in some cases, the device might contain an additional password that is unknown to Security Manager, such as a line console password. If this additional password exists, the device reverts to that password instead of the default password. If that happens, Security Manager cannot configure this device. Therefore, if you use a device access policy to configure the enable password or enable secret password on a device, make sure that you do not unassign the policy without assigning a new policy before the next deployment.
- **Site-to-Site VPNs**—If you lose communication with a spoke in a VPN, the problem can occur when the Security Manager server communicates with an external interface on the spoke from within the hub's protected network. We recommend that when you add the hub device to Security Manager that you define a management IP address that is located outside of the hub's protected network.
- **Platform > Device Admin > Device Access > Allowed Hosts**—For IPS devices, the Allowed Hosts policy identifies the hosts that can connect to the sensor. The Security Manager server must be included in this policy.

Related Topics

- Troubleshooting Device Communication Failures, page 9-8
- Managing Device Communication Settings and Certificates, page 9-4
- Managing IPS Certificates, page 44-10
- Understanding Device Communication Requirements, page 2-1

Updating VPNs That Include Routing Processes

Problem: When you define and deploy changes to a routing process that is being used by a VPN topology (using either the Site-to-Site VPN Manager or the routing policies), the changes that you make are not reflected in the CLI commands configured on the device.

Solution: When you discover a VPN topology that includes routing processes, such as GRE full mesh, Security Manager populates the GRE Modes policy in the Site-to-Site VPN Manager, as well as the relevant routing policies. However, changes made to one of these policies in Security Manager are not automatically reflected in the other policy, which can lead to unexpected results after deployment. Therefore, if you make changes to the secured IGP in the Site-to-Site VPN Manager, be sure to go to Platform > Routing in Device view to make the necessary changes in the device's routing policies. Likewise, if you make changes directly to the routing policy, be sure to make the necessary changes in the Site-to-Site VPN Manager as well.

Related Topics

- Chapter 25, "Managing Site-to-Site VPNs: The Basics"
- Chapter 61, "Managing Routers"
- Chapter 46, "Managing Firewall Devices"

Mixing Deployment Methods with Router and VPN Policies

You might receive unpredictable results when you deploy router platform and VPN policies to a live device after previously deploying to a configuration file.

This problem can occur when you use a mix of deployment methods (deploy to device and deploy to file) with router platform policies and VPN policies. Because Security Manager does not manage all the available CLI commands for these policy types, it maintains a snapshot of the commands it has configured and leaves all other commands (which includes unsupported commands as well as supported commands in policies that have not been configured in Security Manager) intact on the device.

After each deployment, Security Manager creates a snapshot of the policies that were deployed to each device. This snapshot is used during the next deployment to generate the list of configuration changes that will be deployed to the device. Only one snapshot is maintained at a time per device.

Mixing deployment methods with router platform policies and VPN policies can lead to unpredictable results, as shown in this example:

1. Configure router platform policy A to a live device. When deployment completes, Security Manager creates a snapshot for that device with policy A.
2. Next, configure policy B to replace policy A, but instead of deploying policy B to the device, deploy it to a file instead. When this deployment completes, Security Manager creates a snapshot with policy B that replaces the previous snapshot with policy A. However, because you did not deploy policy B to the device, the CLI commands that are required to negate policy A have not been deployed. Policy A is still deployed on the device.
3. Deploy again to the device without first copying the changes in the configuration file to the device. Security Manager cannot generate the commands that are required to negate policy A from the device because the snapshot with policy A no longer exists.

Because policy A is a router platform policy, any of the following results might occur:

- The policy in the latest deployment overrides policy A.
- Both policies end up defined on the device.
- Deployment fails because the two policies cannot coexist.

Therefore, if you deploy to a file when working on a live device, we strongly recommend that you copy your configuration changes from the file to the device before performing additional deployments to the device.

Related Topics

- Chapter 25, “Managing Site-to-Site VPNs: The Basics”
- Chapter 61, “Managing Routers”

Deployment Failures for Routers

Following are some potential problems you might encounter when deploying configurations to Cisco IOS routers.

Deployment Fails for Interface Settings

Problem: Deployment fails for interface settings on a router.

Solution: Security Manager cannot validate whether you have the appropriate types of interface cards or shared port adapters (SPAs) installed on the router, or the appropriate licenses configured, to support your interface policies. If you add or remove an interface card without changing your interface policies,

you can encounter deployment errors. The best practice is to ensure that you discover inventory from the router whenever you change interface modules or SPAs so that Security Manager can discover the appropriate interface features.

Deploying Layer 2 Interface Definitions

Problem: Deployment fails if the interface policy includes a definition for a Layer 2 interface.

Solution: Layer 2 interfaces do not support Layer 3 interface definitions, such as IP addresses. Make sure that you did not define a Layer 3 definition on the Layer 2 interface.

VPN Traffic Sent Unencrypted

Problem: Traffic that should be sent encrypted over a VPN is instead being sent unencrypted.

Solution: Ensure that you are not performing NAT on VPN traffic. Performing address translation on VPN traffic prevents the traffic from being encrypted and sent through the VPN tunnel. When defining dynamic NAT rules, make sure that the Do Not Translate VPN Traffic check box is selected, even when you perform NAT into IPSec. (This option does not interfere with the translation of addresses arriving from overlapping networks.)

This option can be used only on site-to-site VPNs. For remote access VPNs, you need to create an ACL object that explicitly denies the flow containing VPN traffic and define this ACL as part of a dynamic rule in the NAT policy. For more information, see NAT Page: Dynamic Rules, page 24-10.

Unable to Deploy ADSL or PVC Policy

Problem: Deployment fails for your ADSL or PVC policy.

Solution: Make sure that you have selected the correct ATM interface card type in the policy definition. Security Manager cannot properly validate the policy definition without knowing the correct card type, which can lead to deployment failures.

DHCP Traffic Not Being Transmitted

Problem: DHCP traffic is not being transmitted even after you deploy a DHCP policy to the device.

Solution: Check whether an access rule on the device blocks Bootstrap Protocol (BootP) traffic. Having such a rule prevents DHCP traffic from being transmitted.

NAC Not Implemented on Router

Problem: Network admission control is not being implemented on the router, even though a NAC policy was deployed to it.

Solution: Ensure that the default ACL on the router permits UDP traffic over the port defined in the NAC policy for EAP over UDP traffic. This is the protocol that NAC uses for communication between the Cisco Trust Agent (CTA), which is the NAC client that provides posture credentials for the endpoint device on which it is installed and the network access device (NAD; in this case, the router) that relays the posture credentials to the AAA server for validation. The default port used for EAP over UDP traffic is 21862, but you can change this port as part of the NAC policy. If the default ACL blocks UDP traffic, EAP over UDP traffic is likewise blocked, which prevents NAC from taking place.

Deployment Fails with Error Writing to Server or HTTP Response Code 500 Messages

Problem: Deployment to a Cisco IOS router fails and an “Error Writing to Server” or “Http Response Code 500” error message occurs.

Solution: When you use SSL as the transport protocol for deploying configurations to a Cisco IOS router, the configuration is split into multiple configuration bulks. The size of this configuration bulk varies from platform to platform. If Security Manager tries to deploy a configuration bulk that exceeds the size of the SSL chunk configured on that device, the deployment fails and you get an “Error Writing to Server” or “Http Response Code 500” error message.

To resolve this, do the following:

1. On the Security Manager server, open the DCS.properties file in the \CSCOpx\MDC\athena\config folder in the installation directory (usually C:\Program Files).
2. Locate **DCS.IOS.ssl.maxChunkSize=**<*value of the configuration bulk*>.
3. Reduce the value of the configuration bulk.
4. Restart the CiscoWorks Daemon Manager.

Related Topics

- Chapter 61, "Managing Routers"

Deployment Failures for Catalyst Switches and Service Modules

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any enhancements.

Following are some potential problems you might encounter when deploying configurations to Catalyst switches and Catalyst 6500/7600 service modules.

Deployment Fails for Interface Settings

Problem: Deployment fails for interface settings on a Catalyst 6500/7600 device.

Solution: Certain interface settings (such as speed, duplex, and MTU settings) are specific to particular card types and are not validated prior to deployment. Make sure to enter the correct values for your specific card type to ensure successful deployment.

Deployment Failures to FWSM Security Contexts After Changing Interface Policies

Problem: You add an FWSM with security contexts and discover its policies. The configuration includes interface aliases (the allocate interface command). After changing the interfaces policy for a context, deployment fails.

Solution: Connect directly to the FWSM and remove all mapped interface names from the system execution space configuration and in all other contexts, replace interface references to mapped names with the VLAN ID of the interface. You can then delete the FWSM from the Security Manager inventory and rediscover it.

Deployment Failures for FWSMs That Have Multiple Contexts

Problem: Deployment to an FWSM that has multiple security contexts sometimes fails or results in a temporary performance impact to the FWSM.

Solution: The problem is that Security Manager is trying to deploy configurations to more than one security context on a device at the same time. Depending on the configuration changes, this can result in errors on the device that prevent successful deployment. If you use FWSM in multiple-context mode, configure Security Manager to deploy configurations serially to the device so that one context at a time is configured, as described in Changing How Security Manager Deploys Configurations to Multiple-Context FWSM, page 9-17.

Deployment Fails for Internal VLANs

Problem: Deployment fails when Security Manager tries to create a VLAN with an ID that is within the range of the device's internal VLAN list.

Solution: Security Manager cannot detect internal VLANs. Therefore, you must define a VLAN ID that falls outside of the device's internal VLAN list. Use the **show vlan internal usage** command on the device to view the list of internal VLANs.

Deployment Fails When Changing the Running Mode of an IDSM Data Port VLAN

Problem: Deployment fails when you attempt to change the running mode of the data port VLAN from Trunk (IPS) to Capture (IDS) and the following error message is displayed:

Command Rejected: Remove trunk allowed vlan configuration from data port 2 before configuring capture allowed-vlans

Solution: On some software releases such as 12.2(18)SFX4, there is a bug that prevents the change from occurring correctly. Reload the device to overcome the problem.

Deployment Fails for FWSM Configuration With Large Numbers of ACLs

Problem: Deployment to FWSM devices fail when the configuration contains a large number of ACLs.

Solution: This could occur because the CPU utilization is high during ACL compilation. To resolve this, reconfigure the CPU utilization threshold limit by doing the following:

1. On the Security Manager server, open the DCS.properties file in the \CSCOpx\MDC\athena\config folder in the installation directory (usually C:\Program Files).
2. Locate the **DCS.FWSM.checkThreshold=False** property.
3. Change the value to true: **DCS.FWSM.checkThreshold=True**.
4. Restart the CiscoWorks Daemon Manager.
5. Deploy the configuration to the device again.

   After you set the value to true, discovery and deployment checks the CPU utilization and generates error messages if the CPU utilization is not within the configured value set in the DCS.FWSM.minThresholdLimit property. The default value is 85.

Related Topics

- Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers"

Changing How Security Manager Deploys Configurations to Multiple-Context FWSM

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any enhancements.

If you configure a Firewall Services Module (FWSM) to run in multiple context mode, so that you host more than one security context on the FWSM, you need to configure Security Manager to deploy configurations serially to the FWSM. The FWSM has some limitations that can prevent successful deployments if more than one context is updated at the same time, so you might run into deployment failures if you do not use serial deployment. There can also be an impact on FWSM performance during deployment if you do not use serial deployment.

To change how Security Manager deploys configurations to multiple-context FWSM, you need to update the DCS.properties file. You also need to add the FWSM contexts to the inventory using the FWSM admin context, rather than adding the individual security contexts.

The following procedure explains the end-to-end process for ensuring that FWSM deployments are done serially.

Step 1 Make it a standard practice to add FWSM security contexts using the admin context management IP address. Manage the contexts through the admin context.

Although it is possible to add security contexts for an FWSM individually, using each context's management IP address, Security Manager cannot recognize these individually-added contexts as being hosted on the same physical device. This prevents Security Manager from doing serial deployments to the contexts.

If you have any FWSM security contexts that you added using the security context management IP, delete the contexts and FWSM from the inventory, then add them using the admin context (discover all policies). See Adding Devices to the Device Inventory, page 3-7.

Tip If you have any undeployed changes to these contexts that you want to keep, first deploy the changes to ensure that the configurations on the device are complete. Do the deployments one context at a time.

Step 2 Log into Windows on the Security Manager server and edit the **DCS.properties** file in the \CSCOpx\MDC\athena\config folder in the installation directory (usually c:\Program Files). Use a text editor such as NotePad to update the file.

Step 3 Locate the DCS.doSerialAccessForFWSMVCs property in the DCS.properties file and set it to true:

DCS.doSerialAccessForFWSMVCs=true

Step 4 Restart the CiscoWorks Daemon Manager.

Deployment Failures to Devices Managed by AUS

Deployment might fail when deploying to multiple AUS-managed devices after starting the AUS if you perform deployment before the Auto Update Server (AUS) is fully operational. The AUS requires time to start up after the following operations:

- New installation or upgrade.
- Manual restart (including after a power outage).
- Manual restart of the Cisco Security Manager Daemon Manager service.

You can verify whether the AUS is fully operational by verifying the status of its Windows services. To do this, select **Start > Control Panel > Administrative Services > Services**, then check the status of the CiscoWorks AUS Database Engine service. If this service has started, try again to deploy.

Troubleshooting the Setup of Configuration Engine-Managed Devices

The following questions and answers describe issues that might arise when you set up a device managed by a Cisco Configuration Engine (also known as CNS) and how to solve them:

Question: Why does Configuration Engine deployment fail?

Answer: Not all versions of Configuration Engine function in a compatible manner. Because Security Manager does not verify the software version running on a Configuration Engine when you add it to the device inventory, you can add unsupported versions to the inventory. Then, when you try to deploy, you

can run into unpredictable errors. Ensure that you are running a supported version of Configuration Engine (for version information, see the release notes for this version of Security Manager at http://www.cisco.com/en/US/products/ps6498/prod_release_notes_list.html).

Question: Why do I receive an InvalidParameterException when I click on an IOS device on the Configuration Engine web page?

Answer: This is the expected behavior. For IOS devices, Security Manager uses deployment jobs to deploy configurations to Configuration Engine instead of associating a configuration to the IOS device in Configuration Engine. Therefore, you do not see an associated configuration when you click the device name on the Configuration Engine web page. For ASA/PIX devices, Security Manager associates the configuration to the device in Configuration Engine. Therefore, clicking the device name displays the associated configuration.

Question: Why am I getting the following error: com.cisco.netmgmt.ce.websvc.exec.ExecServiceException: [002-01003]]deviceName does not exists?

Answer: This error indicates that the device has not been added to Configuration Engine. It appears if you have not performed rollback or deployment in Security Manager (both of which add the device automatically), and have not manually added the device to Configuration Engine.

Question: Why am I getting the following error: com.cisco.netmgmt.ce.websvc.config.ConfigServiceException: [002-01003]]Device device id is not connected

Answer: The answer depends on the type of setup you are performing:

- Event mode setup—Make sure that the Configuration Engine device ID defined in the Device Properties window in Security Manager matches the device ID configured on the router (using the **cns id string** command).
- Call home mode setup—The device is not connected to Configuration Engine in this mode; therefore, all Security Manager operations that require the retrieval of the device configuration using Configuration Engine are not supported. This includes discovery, preview configuration, display running configuration, and connectivity tests (and rollback, for IOS devices).

Question: Why is deployment to my Configuration Engine-managed ASA/PIX device not working?

Answer: There are several possibilities:

- The configuration contains invalid commands. You can test this by copying the configuration associated with the ASA/PIX device in Configuration Engine and pasting it directly into the device.
- The **auto-update server** command contains an invalid username and password.
- You did not wait long enough for the configuration to be polled into the ASA/PIX device. Use the **show auto** command to verify when the next polling cycle will occur.
- If you previously used the Configuration Engine server for the same ASA/PIX device and did not delete the device from the Configuration Engine server before you started the current task, it is possible that the device received the previous configuration from the server before you deployed the new configuration to it.
- If none of the suggestions above solves the problem, turn on Configuration Engine debug mode on the ASA/PIX device and check the log for errors after the next polling cycle.

Question: Why was I able to deploy successfully to a Configuration Engine-managed ASA/PIX device the first time, but subsequent deployments were unsuccessful?

Answer: This can happen if the configuration pushed during the first deployment contains incorrect CLI commands for the auto-update feature. Check the following:

- Make sure the username and password of the Configuration Engine server is defined correctly in the **auto-update** command.
- If you used **name** commands when configuring the auto-update server using the device CLI, make sure that you have defined a FlexConfig that contains the necessary **name** commands. A FlexConfig is necessary because Security Manager does not support this command directly. As a result, even though the command was discovered, it does not appear in the full configuration. If you use Security Manager to configure the AUS policy, **name** commands are not necessary.

Question: How do I debug Configuration Engine on an ASA/PIX device?

Answer: Enter the following CLI commands:

```
logging monitor debug
terminal monitor
logging on
```

You can also find relevant information in the PIX log on the Configuration Engine server.

Question: How do I debug Configuration Engine on an IOS device?

Answer: Enter the following CLI commands:

```
debug cns all
debug kron exec-cli
terminal monitor
```

When working in event mode, you can also find relevant information in the event log on the Configuration Engine server. When working in call home mode, check the config server log on the Configuration Engine server.

Question: Why did I fail to discover an IOS device and acquire its configuration through Configuration Engine?

Answer: If you see the following errors in debug mode:

*Feb 23 21:42:15.677: CNS exec decode: Unknown hostname cnsServer-lnx.cisco.com ... 474F6860: 72726F72 2D6D6573 73616765 3E584D4C error-message>XML 474F6870: 5F504152 53455F45 52524F52 3C2F6572 _PARSE_ERROR

Verify the following:

- The CNS commands use a fully-qualified host name (host name and domain name).
- The device contains the **ip domain name** command.
- The device contains the **ip host** command with the fully qualified hostname of the Configuration Engine with its IP address.

Question: Why does the event mode router not appear on the Configuration Engine Discover Device page or appear in green on the Configuration Engine web page?

Answer: Check the following:

- Make sure that the router and the Configuration Engine server can ping each other.
- Make sure that the event gateway on the Configuration Engine server is up and running by using one of the following commands:

  Status for plain text mode: **/etc/init.d/EvtGateway**

  Status for SSL encrypted mode: **/etc/init.d/EvtGatewayCrypto**
- Clear the **cns event** command, then re-enter it without specifying a port number.

Managing the Security Manager Server

The following topics describe some system management tasks related to the general operation of the Security Manager product:

- Overview of Security Manager Server Management and Administration, page 10-1
- Managing a Cluster of Security Manager Servers, page 10-2
- Installing Security Manager License Files, page 10-16
- Certificate Trust Management, page 10-18
- Working with Audit Reports, page 10-19
- Taking Over Another User's Work, page 10-23
- Changing Passwords for the Admin or Other Users, page 10-24
- Backing up and Restoring the Security Manager Database, page 10-24
- Generating Data for the Cisco Technical Assistance Center, page 10-28

Overview of Security Manager Server Management and Administration

As a software application, Cisco Security Manager runs on the framework provided by the CiscoWorks Common Services application. Many of the fundamental server control functions are provided by Common Services. For example, if you want to create a multiple-server setup for Security Manager, you must create that setup in Common Services. Common Services also provides the tools for creating and managing local user accounts, for backing up and restoring the database, for generating various reports on system functions, and for many other basic functions.

To access the Common Services application, do any of the following:

- If you currently have the Security Manager client open, you can select **Tools > Security Manager Administration** and select **Server Security** from the table of contents. The Server Security page has buttons that link to and open specific pages in Common Services. You can click any button and then navigate to any desired page in Common Services.
- Using your web browser, link to the Security Manager server using the URL https://*servername*, where *servername* is the IP address or DNS name of the server. This URL opens the Security Manager home page. Click **Server Administration**, or the **CiscoWorks** link, to open Common Services.

Note A special consideration applies if you are using Internet Explorer 10.x in Windows Server 2012 (Standard or Datacenter)—64-bit, support for which is new in Version 4.7 of Cisco Security Manager. You need to be aware of this consideration when using the following navigation path: Windows Start > Cisco Security Manager Client > [log in] > Configuration Manager > Tools > Security Manager Administration... > Server Security. The Server Security page will open normally for you, but on that page you will be unable to use the buttons (e.g., Local User Setup) to cross-launch the Server Security Tools within Common Services. To work around this problem, decrease the security levels of your intranet settings in Internet Explorer 10.x.

To learn more about the things you can do with Common Services, browse the Common Services online help.

Note The **Software Center > Software Update** feature in Common Services is not supported by Cisco Security Manager.

Managing a Cluster of Security Manager Servers

A Security Manager server cluster is two or more Security Manager servers used to manage a network. Typically, you want to maintain some relationship between the servers. Although there is no systematic relationship between the servers in the cluster, there are some techniques that you can use to maintain a cluster-like relationship. The topics in this section explain how you can manage a group of Security Manager servers as a cluster.

This section contains the following topics:

- Overview of Security Manager Server Cluster Management, page 10-2
- Exporting the Device Inventory, page 10-6
- Exporting Shared Policies, page 10-12
- Importing Policies or Devices, page 10-13

Overview of Security Manager Server Cluster Management

You can manage a large number of devices with a single Security Manager server. There are, however, a variety of reasons for managing your network with more than one Security Manager server. For example:

- If you have a very large network with thousands of devices to manage, you might find performance to be unacceptable when trying to manage all devices from a single server.
- For geographic reasons, you might find it better to have servers that are closer to managed devices. For example, if you have major sites on different sides of the globe, having separate servers at each major site might simplify management and improve performance. For example, when deploying configurations to managed devices, a Security Manager server located in Bangalore should be able to deploy configurations to a device in Bangalore much faster than a Security Manager server located in San Francisco simply due to the much shorter physical network distance.

- You might want to segment device management based on the technology managed. For example, you might want to use one server to manage your site-to-site VPNs, another server to manage ASA firewall and remote access VPN policies, and a third server to manage IPS.
- Separate IT organizations might be managing different parts of your network. Although you can set up ACS to fine-tune access control to the device level, you might instead find it simpler to have distinct Security Manager servers for each IT organization.

If you decide to install more than one Security Manager server, the main challenges are the following:

- Splitting a single server into two or more servers—You might currently have a single Security Manager server, and decide that you need multiple servers. For information on how to split a Security Manager server into two or more servers, see Splitting a Security Manager Server, page 10-3.
- Maintaining the same set of shared policies—If you use multiple servers to manage the same device types, you might want to ensure that the shared policies assigned to the devices are identical. For example, you might want to have the same set of mandatory and default access rules inherited by all ASA devices.

  There is no automatic process for maintaining the same set of shared policies among a cluster of servers. Instead, you must manually export them from your main server and import them into the remaining servers. For more information, see Synchronizing Shared Policies Among Security Manager Servers, page 10-5.

Splitting a Security Manager Server

If you decide that you need to convert a single Security Manager server into two or more servers, you can split the server by moving subsets of the devices managed by the original server to the new servers. Keep in mind that you should manage a specific network device from a single Security Manager server, so delete the moved devices from the original server.

Tip Use the same release of Security Manager software on all servers.

Related Topics

- Overview of Security Manager Server Cluster Management, page 10-2
- Synchronizing Shared Policies Among Security Manager Servers, page 10-5
- Exporting Shared Policies, page 10-12

Step 1 Install the new Security Manager servers as described in the *Installation Guide for Cisco Security Manager*.

Ensure that the server is functioning correctly, and also ensure that you install licenses with a device count that will be sufficient for the devices you will move to the server. Ensure that you use a professional license if you manage device types that require it. For information on installing licenses, see Installing Security Manager License Files, page 10-16.

Step 2 On the original server, verify that the policies of the devices that you will move will allow access from the IP address of the new server. For example, consider access rules on ASAs and routers, and the Allowed Hosts policy on IPS devices.

Step 3 On the original server, ensure that all configuration changes for the devices you are moving have been submitted and deployed. You will need to ask the staff to submit and deploy their changes, there is no simple way to determine this status within Security Manager.

This step ensures that there are no pending uncommited changes. For information on deploying configurations, see the following topics based on workflow mode:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Step 4 Select **File > Export > Devices** to export the devices with their assigned policies and policy objects from the original Security Manager server. Be sure to select **Export Devices, Policies, and Objects** during the device export so that policy information is included. The file type must be **dev**. For more detailed information, see Exporting the Device Inventory from the Security Manager Client, page 10-6.

Create separate export files containing unique devices for each new Security Manager server.

Tip At this point, do not make policy changes to the exported devices in the original server, and do not deploy configurations to those devices. If you find that you need to make changes to the devices from the original server before you complete the split, create a new export file.

Step 5 On each of the new Security Manager servers, select **File > Import** to import the exported information to the new servers. For more detailed information, see Importing Policies or Devices, page 10-13.

Tip Device groups are not preserved during import. All devices are placed in the All group. You need to manually recreate the desired device group structure and add the devices to the appropriate groups.

Step 6 Verify that each of the new Security Manager servers can manage the newly-imported devices. For example, you could do a deployment, even for unchanged devices, to ensure that the new server can successfully contact all devices and deploy configurations.

Tip As explained in Importing Policies or Devices, page 10-13, you must submit policies before the changes are available for configuring devices. Submit policies before doing a deployment.

Step 7 If you were monitoring any of the moved devices using the original server (that is, with Event Viewer and optionally Report Manager), ensure that you update the relevant policies to have syslog messages sent to the new server and to allow contact from the new server. None of the event or report data from the original server is transferred to the new server.

For information on configuring the devices to enable Security Manager monitoring, see the following topics:

- Configuring ASA and FWSM Devices for Event Management, page 69-28
- Configuring IPS Devices for Event Management, page 69-29

Step 8 On the original Security Manager server, select **File > Delete Devices** to delete the moved devices from the original server. For information on deleting devices, see Deleting Devices from the Security Manager Inventory, page 3-58.

Synchronizing Shared Policies Among Security Manager Servers

When you have more than one Security Manager server, you can manually synchronize the shared policies among those servers. When you synchronize shared policies, the policy objects that are used by those shared policies are also synchronized.

Tips

- There is no programmatic way to identify a single Security Manager server as the "primary" server, the one that contains the official version of shared policies. You must decide which server to use as the primary and have the discipline to edit shared policies on that server only.
- Use the same release of Security Manager software on all servers.
- You can also synchronize certain types of policy object among servers even if those objects are not used in shared policies. If you have network/host, service, or port list objects that you want to synchronize, you can use the command described in Importing and Exporting Policy Objects, page 6-23.
- When importing shared policies and policy objects, the imported information always replaces any existing shared policies or policy objects of the same name. Therefore, if you allow users to create their own shared policies and objects on a server where you will import policies and objects, it is critical that you develop a policy and object naming standard so that user policies and objects are not accidentally overwritten by newly imported policies and objects.

Related Topics

- Overview of Security Manager Server Cluster Management, page 10-2
- Splitting a Security Manager Server, page 10-3
- Exporting the Device Inventory from the Security Manager Client, page 10-6

Step 1 On the original server, ensure that all configuration changes for the shared policies and policy objects have been submitted. You will need to ask the staff to submit their changes and have them approved, there is no simple way to determine this status within Security Manager.

When exporting shared policies, there is no need to ensure that new changes have been deployed to devices assigned to the policies. Device assignments and deployment status are not part of the exported information.

Step 2 Select **File > Export > Policies** to export the shared policies and any policy objects used by the policies. The export process creates a file with the extension **pol**.

Tip You cannot pick and choose which policies to export. You can select policy types only. All shared policies of a selected type are exported.

For more detailed information, see Exporting Shared Policies, page 10-12.

Step 3 On each of the other Security Manager servers, select **File > Import** to import the exported shared policy information to the servers. For more detailed information, see Importing Policies or Devices, page 10-13.

Tip Any shared policies or objects that have the same name as imported ones are replaced. The import of a policy or object will fail if a user already has a lock on the policy or object. As explained in Importing Policies or Devices, page 10-13, you must submit policies before the changes are available for configuring devices.

Step 4 If you do not want to import all of the shared policies, delete the ones you did not want to import on the other servers. This is a manual process.

Exporting the Device Inventory

Exporting the device inventory allows you to import the inventory into other network management applications, or to manipulate the output for your own reporting purposes. There are two unrelated methods to export the device inventory:

- Use the **File > Export > Devices** command—Using this command, you can create either a simple comma-separated values (CSV) file, or a compressed .dev file that contains the devices along with their complete configuration policies. The CSV file is in a format suitable for importing into the CiscoWorks Common Services Device Credential Repository (DCR), the Cisco Security Monitoring, Analysis and Response System (CS-MARS), Cisco Prime Security Manager (PRSM), or another Security Manager installation; or you can open it and view it in a spreadsheet or text editor program. The .dev file is suitable for importing into another Security Manager server only. For more information, see Exporting the Device Inventory from the Security Manager Client, page 10-6.
- Use the CSMgrDeviceExport Perl script—Using this Perl script, you can export the inventory without starting the Security Manager client. You can direct the output to the screen or to a comma-separated values (CSV) file. For more information, see Exporting the Device Inventory from the Command Line, page 10-10.

Note Only the five most recent versions of each device configuration are exported.

Exporting the Device Inventory from the Security Manager Client

You can export the device inventory in a variety of formats. These are the main choices:

- **Export as CSV** (comma-separated values)—You can create a simple CSV file containing inventory information in one of the following formats: CSM (for use with Cisco Security Manager), Device Credential Repository (DCR, for CiscoWorks Common Services), and CS-MARS seed file (for use with Cisco Security Monitoring, Analysis and Response System). You can open a CSV file in a spreadsheet application or text editor, and use the file with any application that supports the format, including other Security Manager servers. However, this format contains no policy information, so if you use it with another Security Manager server, you must discover policies while adding the devices.
  - For more information about the CSV formats, see Supported CSV Formats for Inventory Import/Export, page 10-9.
  - For information on how to import devices from a CSV file, see Adding Devices from an Inventory File, page 3-31.

- **Export Devices, Policies, and Objects**—Export the device inventory along with all device properties, policies, and policy objects used by the device. Exported information includes the following:

  Note Importing of *.pol or *.dev files is only supported on the same version of Cisco Security Manager as used when exporting those files. You cannot export from one version of Cisco Security Manager and import on a server running a different version.

  - All local and shared policies assigned to the devices, including all policy objects used in the policies and any device-level overrides for the objects. Shared policy assignments are maintained.
  - Device properties and inventory.
  - Configuration Archive data for the devices.
  - History snapshots for the devices.
  - Device certificates.
  - IPS device license and certificate information. Applied signatures are not exported (when importing the device, you must have the same signature package registered on the server). IPS update settings are not included; you will have to recreate them after import.
  - The VPN topologies in which the devices participate. However, a VPN topology is exported only if all devices that participate in the topology are included in the export. Extranet VPNs are always exported.

  Thus, the export file includes the complete policy configuration for the selected devices. The file created has the extension .dev and can be read only by another Security Manager server (the file contents are compressed and uninterpretable, which preserves the security of your policy information).

  For information on importing a .dev file into another Security Manager server, see Importing Policies or Devices, page 10-13.

Export Size Limitations

If your Security Manager database contains a large number of devices or a large number of policies or policy objects, you should limit the number of devices you export at one time to prevent errors. The following guidelines can be used to help estimate the number of devices you can successfully export at one time:

Example 1: 1000+ devices in the database with approximately 1500+ polices per device and approximately 25,000 objects in the database:

- Maximum number of devices (devices only) to be exported at one time = 250
- Maximum number of devices (along with polices and objects) to be exported at one time = 100 to 150

Example 2: Fewer than 1000 devices in the database with approximately 1500+ polices per device and approximately 10,000-15,000 objects in the database:

- Maximum number of devices (devices only) to be exported at one time = 250 to 300
- Maximum number of devices (along with polices and objects) to be exported at one time = 200

Tips

- When you select the **Export Devices, Policies, and Objects** option, you can export to the Security Manager server or to the local Security Manager client. When exporting a CSV file, you can only export to the Security Manager server. You can control the ability to export to or import from the local Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10.
- Exported devices are not deleted from the inventory. If you intend to manage the devices from a different Security Manager server, delete the devices after successfully importing them into the other server.
- If you select a device that uses an AUS or Configuration Engine to manage its configuration, you should also select the server in the list of devices to export. You cannot export AUS or Configuration Engine information in CS-MARS format.
- You can export unmanaged devices.
- When exporting devices with their policies, only policies and policy objects that have been submitted and approved are included in the export file. Make sure that all desired submissions and approvals have occurred before exporting devices with policies and policy objects.
- No type of export file includes event and report data (that is, data that is available through Event Viewer or Report Manager). Thus, if you are exporting devices with the intention of moving them to another Security Manager server, the event and report data that was already collected for the device will not be available on the new server.
- No type of export file includes device group information. You will have to manually recreate device groups and assign devices to them after importing devices.
- When selecting security contexts or virtual sensors, be sure to select the host device as well. Also, if a device is part of a VPN, ensure that you select all devices in the VPN when exporting devices, policies, and policy objects.
- When selecting IPS or IOS IPS devices, make sure that you have already applied an IPS signature update to the device. Although you can export an IPS or IOS IPS device with the base sensor package (Sig0), you will not be able to import it. The import error will be “missing Sig0 package.”
- When you select the following types of device that are contained in another device, the hosting device is also automatically exported: any module in a Catalyst 6500/7600, an AIM or NME module in a router. You can separately export ASA devices and their IPS modules.
- You cannot export devices with their policies (.dev format) while an activity or configuration session is being approved. All approvals must be complete before you can export devices with policies. In Workflow mode with an approver, contact the approver and ask that the approval be completed in a timely manner. In non-Workflow mode, or Workflow mode without an approver, wait a few minutes before retrying the export, because approvals happen automatically when changes are submitted.
- While you are exporting devices with their policies (.dev format), policy changes cannot be approved. Once the export file has been created, and the command finishes, users can again approve policy changes. In non-Workflow mode, or Workflow mode without an approver, this means that submissions are not allowed during the export process.
- To export devices, policies, and policy objects, you must have Modify Policy and Modify Object privileges to the policy and object types, and Modify Device privileges. These privileges can be assigned for separate policies, objects, and devices when using ACS for authorization control. Having system administrator, network administrator, or security administrator privileges provide the required privileges.

  When exporting devices to CSV, you need Modify Devices privileges only.

Related Topics

- Filtering Items in Selectors, page 1-45
- Selecting or Specifying a File or Directory in Security Manager, page 1-50
- Customize Desktop Page, page 11-10

Step 1 In Device view, select **File > Export > Devices** to open the Export Inventory dialog box.

Step 2 Select either **Export as CSV** or **Export Devices, Policies, and Objects**. These options are described above.

Step 3 Select the devices that you want to include in the export file and click **>>** to add them to the Selected Devices list. You can select a folder to select all devices in the folder.

The list from which you choose contains only those devices to which you have Modify Device permissions.

Note If your Security Manager database contains a large number of devices or a large number of policies or policy objects, you should limit the number of devices you export at one time to prevent errors. For more information, see **Export Size Limitations** above.

Step 4 Click **Browse** to select the folder in which to create the export file and to enter a name for the file. For File Type, select the type of file that you want to create; this selection is critical when creating a CSV file, whereas there is a single option when creating a .dev file.

Click **Save** to return to the Export Inventory dialog box. The Export Inventory To field is updated with the export file information.

Step 5 Click **OK** to create the export file.

A message indicates when the export completes and whether there were errors in the export. When you click **OK**, if there are problems during the export, a dialog box opens listing the messages. If there is a Details button in the dialog box, you can select a message and click **Details** to see the message in a more readable format.

Supported CSV Formats for Inventory Import/Export

When you export devices to a CSV (comma-separated values) file (by selecting **File > Export > Devices** and selecting Export as CSV), or import devices from a CSV file (by selecting **File > New Device** and then selecting Add Device from File in the New Device wizard), you can select one of the following CSV file formats:

- **Device Credential Repository (DCR)**—The device management system for CiscoWorks Common Services. For information on this format, see the description of the sample version 3.0 CSV file in the Common Services documentation at this URL: http://www.cisco.com/en/US/docs/net_mgmt/ciscoworks_common_services_software/3.3/user/guide/dcr.html#wp1193611
- **CS-MARS seed file**—Cisco Security Monitoring, Analysis and Response System. For information on this format, see the CS-MARS documentation at this URL: http://www.cisco.com/en/US/docs/security/security_management/cs-mars/6.0/device/configuration/guide/chDvcOver.html#wp162016

- **Cisco Security Manager**—The Security Manager format, which is the DCR version 3.0 format with additional fields. If you are importing the inventory into another Security Manager server, selecting this format will allow you to import the inventory without discovering policies on the devices.

  Note If the file does not specify os_type and os_version for a device, you must discover policies directly from the device when adding it.

  The additional fields, which appear at the end of each row, are:

  - os_type. The operating system type, which can be one of the following: PIX, ASA, IOS, FWSM, IPS. This field is required for all device types.
  - os_version. The target operating system version, which can be any of the version numbers listed in the New Device wizard when you select Add New Device. The acceptable version numbers differ depending on the device model, so if you are creating a CSV file by hand, look over this list carefully. For more information about adding devices using this method, see Adding Devices by Manual Definition, page 3-26. This field is required for all device types.
  - fw_os_mode. The mode in which a firewall device is running, which can be one of the following: TRANSPARENT, ROUTER, MIXED. This field is required for ASA, PIX, and FWSM devices.
  - fw_os_context. The context in which a firewall device is running, which can be one of the following: SINGLE, MULTI. This field is required for ASA, PIX, and FWSM devices.
  - anc_os_type. The ancillary operating system type for Cisco IOS-IPS devices. If present, it is IPS. This field is required for IOS IPS devices.
  - anc_os_version. The ancillary target operating system version, which is the IPS target operating system version. If present, it can be any of the supported IOS-IPS versions. This field is required for IOS IPS devices.

You can use these CSV files with any program that supports the file format. You can also create a CSV file yourself and use the file to import inventory into Security Manager.

Related Topics

- Exporting the Device Inventory from the Security Manager Client, page 10-6
- Importing Policies or Devices, page 10-13
- Adding Devices from an Inventory File, page 3-31

Exporting the Device Inventory from the Command Line

Security Manager includes a Perl script that you can use to export the device inventory without starting the Security Manager client. You can use this script to automate various offline reporting tasks that your organization might require. You can pipe the output to a comma-separated values (CSV) file or otherwise capture and manipulate the output.

Tip This command does not produce a file that you can use for importing devices or for adding devices "from file." Although it exports inventory information, making it similar to the integrated export features, the usefulness of the command is limited to reporting purposes for organizations that have unique off-line reporting process requirements.

The Perl command is located in $NMSROOT\bin, which is typically C:\Program Files\CSCSpx\bin. The syntax of the command is:

perl [*path*]**CSMgrDeviceExport.pl -u** *username* [**-p** *password*] [**-s** {**Dhdoirtg**}] [**-h**] [**>** *filename.csv*]

Syntax

perl [*path*] **CSMgrDeviceExport.pl**	The Perl script command. Include the path to the CSMgrDeviceExport.pl file if the path is not defined in the system path variable.
-u *username*	A Security Manager username. The data exported is limited by the permissions assigned to this user. The user must have View Device permissions.
-p *password*	(Optional.) The user's password. If you do not include the password on the command, you are prompted for it.
-s {**Dhdoirtg**}	(Optional.) The fields you are selecting to include in the output. If you do not specify the -s option, all fields are included. You can specify one or more of the following: • D—Display name. • h—Host name. • d—Domain name. • o—Operating system (OS) type. • I—Image name. • r—Running OS version. • t—Target OS version. • g—Device groups.
-h	(Optional.) Display the command line help. If you include this option, all other options are ignored.
> *filename.csv*	(Optional.) Pipe the output to the specified file. If you do not specify a file, the output is displayed on the screen.

Output Format

The output is in standard comma-separated values (CSV) format, which you can open in spreadsheet programs or process with your own scripts. The first line has column headings. The columns, left to right, are in the order of the fields described for the -s option above.

If there is no value for a particular field, that field is blank in the output.

The device group output field is enclosed in double-quotes and it can contain more than one group name. The group names include the path structure for the group. For example, the following output indicates the device is part of two groups, the East Coast group in the Department folder, and the NewGroup group in the New folder. Groups are separated by a semicolon.

```
"/Department/East Coast; /New/NewGroup"
```

Any error messages generated during the script are written to the output file.

Exporting Shared Policies

You can export shared policies and the policy objects that they use so that you can import them into another Security Manager server. This can help you maintain the same policies among a group of servers, as explained in Synchronizing Shared Policies Among Security Manager Servers, page 10-5.

Note Importing of *.pol or *.dev files is only supported on the same version of Cisco Security Manager as used when exporting those files. You cannot export from one version of Cisco Security Manager and import on a server running a different version.

Tips

- You can export to the Security Manager server or to the local Security Manager client. You can control the ability to export to or import from the local Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10.
- Only shared policies and policy objects that have been submitted and approved are included in the export file. Make sure that all desired submissions and approvals have occurred before exporting policies.
- All policy objects referenced by shared policies are also exported. However, if a policy object is not referenced, it is not exported. There is a separate command you can use to export network/host, service, and port list objects directly; for information, see Importing and Exporting Policy Objects, page 6-23.
- You cannot export policies while an activity or configuration session is being approved. All approvals must be complete before you can export policies. In Workflow mode with an approver, contact the approver and ask that the approval be completed in a timely manner. In non-Workflow mode, or Workflow mode without an approver, wait a few minutes before retrying the export, because approvals happen automatically when changes are submitted.
- While you are exporting policies, policy changes cannot be approved. Once the export file has been created, and the command finishes, users can again approve policy changes. In non-Workflow mode, or Workflow mode without an approver, this means that submissions are not allowed during the export process.
- To export policies and their policy objects, you must have Modify Policy and Modify Object privileges to the policy and object types. These privileges can be assigned for separate policies, objects, and devices when using ACS for authorization control. Having system administrator, network administrator, or security administrator privileges provide the required privileges.

Related Topics

- Overview of Security Manager Server Cluster Management, page 10-2
- Splitting a Security Manager Server, page 10-3
- Synchronizing Shared Policies Among Security Manager Servers, page 10-5
- Exporting the Device Inventory from the Security Manager Client, page 10-6
- Selecting or Specifying a File or Directory in Security Manager, page 1-50
- Customize Desktop Page, page 11-10

Step 1 In Configuration Manager, select **File > Export > Policies** to open the Export Shared Policies dialog box.

Before opening the dialog box, Security Manager evaluates and loads any shared policies that are defined.

Step 2 Select the shared policies that you want to export, using any of the following methods:

Tip You can use multiple methods to select policies. For example, you can select all shared policies modified since a certain date and also select all shared policies of a specific type to export those policies in the same file.

- To select all shared policies that have been modified since a certain date, enter that date in the Modified since field and click **Select >>** next to the Modified since field. You can enter the date in *MMM DD YYYY* format or you can click the Calendar to select the desired date.
- To select all shared policies, select the **All** folder and click **Select >>** under Browse All Shared Policies.
- To select all shared policies of a specific type, select the shared policy type and click **Select >>** under Browse All Shared Policies. You can select a folder to move all types within the folder to the selected list.
- To specify specific shared policies to export, select the type of shared policies you want to export from the Browse All Shared Policies list, select the check boxes next to the shared policies of that type that you want to export, and then click **Select >>** to move them to the Selected Policies list. If you do not select any specific shared policies, all policies of the selected type will be added to the Selected Policies list.

Note Only those policy types for which shared policies have been defined are listed.

To remove any policies from the Selected Policies list, select them and then click the **<< Remove** button. Use the **Select All** checkbox to specify that all entries in the Selected Policies list are to be removed.

Step 3 Click **Browse** next to the Export Shared Policies To field to select the folder where the export file is to be created, and enter a name for the file. The file type is pre-selected as .pol; you cannot change the file type.

Click **OK** to save the file name and location.

Step 4 Click **OK** in the Export Shared Policies dialog box to begin the export. When the export is completed, you are told how many shared policies were exported and if there are warnings or errors, a dialog box opens listing the problems.

You can now import the policies into other Security Manager servers as explained in Importing Policies or Devices, page 10-13.

Importing Policies or Devices

You can import a shared policy (.pol) or device inventory plus policies (.dev) file that was exported from another Security Manager server.

Note Importing of *.pol or *.dev files is only supported on the same version of Cisco Security Manager as used when exporting those files. You cannot export from one version of Cisco Security Manager and import on a server running a different version.

Tips

- You can import from the Security Manager server or from the local Security Manager client. You can control the ability to export to or import from the local Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10.
- When importing devices, the server must have a sufficient Security Manager license to support the number and types of devices that you are importing. Ensure that you install a professional license before importing device types that require it. For information on installing licenses, see Installing Security Manager License Files, page 10-16.
- When importing policies, whether during device or shared policy import, only the policy types selected for management on the Security Manager Administration Policy Management page will be visible. However, all policies are imported. If you select a previously deselected policy type for management, those imported policies appear with their imported configurations. For more information about selective policy management, see Customizing Policy Management for Routers and Firewall Devices, page 5-11.
- When importing shared polices and policy objects, if a policy or object on the server has the same name as an imported one, it is replaced by the imported policy or object. If there are locks on the policy or object, the import for that policy or object will fail. The message will indicate that the failure was due to a locking problem. To avoid problems, ensure that all users have submitted and approved any changes to shared policies or policy objects before doing an import.
- When importing devices, any shared policies and policy objects assigned to the device are also imported, and these policies and objects replace existing policies and objects under the same conditions as used when importing shared policies.
- To import policies and their policy objects, you must have Modify Policy and Modify Object privileges to the policy and object types. When importing devices, you must also have Modify Device privileges. These privileges can be assigned for separate policies, objects, and devices when using ACS for authorization control. Having system administrator, network administrator, or security administrator privileges provide the required privileges.
- You can import a file only if it was exported from a server running the same release of Security Manager.
- You cannot import a device if the device is already in the inventory. Thus, you cannot update device policies from an import file. If you want to re-import a device, first delete it from the inventory.
- When importing devices that use AUS or Configuration Engine servers to manage configuration deployment, the servers must either be included in the import file or already defined in the Security Manager server, but not both. You will get duplicate display name errors if the import file includes an AUS or Configuration Engine already defined in the inventory. You will get an “invalid server selection” error if you try to import a device that has an AUS or Configuration Engine server assigned to it, but the server is not included in the import file or defined in the inventory.
- You can import unmanaged devices.

- When importing IPS devices, the server must have the same signature levels as the imported devices. For example, if you import two IPS devices, one running signature level 481 and the other 530, you must have both 481 and 530 installed on the server. You might need to download signature packages before importing IPS devices as described in Checking for IPS Updates and Downloading Them, page 44-5.
- This procedure explains how to import .pol or .dev files. If you want to import a device inventory from a CSV file, the procedure is explained in Adding Devices from an Inventory File, page 3-31. The procedures are not similar.

Related Topics

- Overview of Security Manager Server Cluster Management, page 10-2
- Splitting a Security Manager Server, page 10-3
- Synchronizing Shared Policies Among Security Manager Servers, page 10-5
- Selecting or Specifying a File or Directory in Security Manager, page 1-50
- Customize Desktop Page, page 11-10

Step 1 In Configuration Manager, select **File > Import** to open the Import dialog box.

Step 2 Click **Browse** to select the file. Make sure that you select the desired file type (either .pol or .dev) from the Files of Type list on the Select a File dialog box.

Click **OK** when you have selected the file.

Step 3 Click **OK** in the Import dialog box.

You are warned that imported policies or policy objects will replace same-named policies and objects. If you have the required authorization privileges (system administrator or Modify Admin), you have the option to deselect **Display a warning on all shared policies and imported objects**. When selected, the banner for shared policies and imported objects warns users that shared policies might have been created during import, and that specific objects were in fact created during import. This warning provides notice that if the user changes the policy or object, those changes might be overridden by a subsequent policy import. Select whether you want to display a warning and click **Yes**.

Tip If you decide later that you want to change whether the warning is displayed, you can modify the **Display a warning on all shared policies and imported objects** option on the **Tools > Security Manager Administration > Policy Management** page.

The information is imported and you are informed of the results. If errors occur, nothing is imported and a dialog box opens that explains the errors. The most common errors include duplicate device display names when importing devices, or locks on shared policies or policy objects that have the same name as those being imported.

- To resolve the duplicate display name problem, you must delete the device from the inventory or rename it. You cannot selectively import devices, you must import all or none.

Note Not all duplicate device names might be listed. When using AUS or Configuration Engine to manage configuration deployment, imported AUS and Configuration names are evaluated before managed device names. Thus, you might see new duplicate display name errors after fixing the first set of errors.

- To resolve the locking problem, you must ensure that users submit their policy changes and have them approved. When importing devices, you might have to delete the imported devices before retrying the import.

Tip When importing devices, it might take some time for the devices to appear in the device list in Device view. Also, device groups are not preserved during import. All devices are placed in the All group. You need to manually recreate the desired device group structure and add the devices to the appropriate groups.

Step 4 Because policy changes are performed under an activity or configuration session, the imported policies and policy objects are not yet committed to the Security Manager database. You must submit and approve the changes. Based on Workflow mode:

- Non-Workflow mode—Select **File > Submit**.
- Workflow mode without an approver—Select **Activities > Approve Activity**.
- Workflow mode with an approver—Select **Activities > Submit Activity**. The activity must be approved before the changes are committed.

If you are not happy with the import, you can discard the activity or configuration session. However, when importing devices, the devices are added outside an activity or configuration session. Therefore, if you discard the activity or configuration session, you discard the device policies and VPN topologies, but the devices remain in the inventory. You should also delete the devices as described in Deleting Devices from the Security Manager Inventory, page 3-58.

Installing Security Manager License Files

The terms of your Security Manager software license determine many things, including the features that are available to you and the number of devices that you can manage. For licensing purposes, the device count includes any physical device, security context, virtual sensor, or Catalyst security services module that uses an IP address. Failover pairs count as one device. For PIX Firewalls, FWSM, and ASA devices that are configured in multiple-context mode (so that they host more than one security context), only the security contexts are counted as devices; the hosting device is not counted as a separate device.

Three license types, Standard, Professional, and Upgrade, are available, in addition to a free 90-day evaluation period that is restricted to 50 devices. For complete information on the types of licenses available and the various supported upgrade paths, as well as information about the Cisco Software Application Support service agreement contracts that you can purchase, see the product bulletin for this version of Security Manager at http://www.cisco.com/en/US/products/ps6498/prod_bulletins_list.html. Also see the *Installation Guide for Cisco Security Manager*.

License limits are imposed when you exceed the allotted time (in the case of the evaluation license), or the number of devices that your license allows you to manage. The evaluation license provides the same privileges as the Professional Edition license. It is important that you register Security Manager as soon as you can within the first 90 days, and for the number of devices that you need, to ensure uninterrupted use of the product. Each time you start the application you are reminded of how many days remain on your evaluation license, and you are prompted to upgrade during the evaluation period. At the end of the evaluation period, you are prevented from logging in until you upgrade your license.

For non-evaluation licenses, if the database contains more devices than what is allowed by your configured licenses, you cannot log into the application using the Security Manager client. You are prompted to add a license during login, and login cannot complete until you add an appropriate license.

Tip The number of devices includes all discovered security contexts and virtual sensors, even if you have not submitted the activity that discovered them and they do not currently appear in the device selector. If it appears there are fewer devices in the inventory than your license allows, but you are getting device count error messages, submit all activities to determine the number of discovered devices. Delete those that you do not want to manage.

Before You Begin

- Obtain the base or upgrade license and any additional licenses that you require. You must have a Cisco.com user ID, and you must register your copy of the software on Cisco.com. When registering, you must provide the Product Authorization Key (PAK) that is attached to the Software License Claim Certificate inside the shipped software package.
  - If you are a registered Cisco.com user, go to http://www.cisco.com/go/license.
  - If you are not a registered Cisco.com user, go to http://tools.cisco.com/RPF/register/register.do.

  After registration, the base software license is sent to the e-mail address that you provided during registration. In addition to receiving a PAK and license for Security Manager, you might receive one additional PAK for each incremental device count pack you purchased.

  Copy the license files to a folder on the Security Manager server or your local Security Manager client. If you are copying the license files to your Security Manager server, you must store your license files on a disk that is local to your Security Manager server; you cannot use a drive that is mapped to the server. Windows imposes this limitation, which serves to improve Security Manager performance and security.

  Note To install a license file that is located on your local Security Manager client, you must have client-side file browsing enabled (see Customize Desktop Page, page 11-10).

Tip Do not place the license file in the etc/licenses/CSM folder in the product's installation folder on the Security Manager server, or you will encounter an error when you try to add the license. Place the file in a folder outside the product folders.

- Common Services does not require a license file.
- Auto Update Server does not require a license file.

Step 1 Select **Tools > Security Manager Administration** and select **Licensing** from the table of contents.

Step 2 Click **CSM** if the tab is not active. For a description of the fields on this tab, see CSM Tab, Licensing Page, page 11-57.

Step 3 Click **Install a License** to open the Install a License dialog box.

The Install a License dialog box includes links to Cisco.com for obtaining licenses if you have not already done so. Click **Browse** to select a license file, and then click **OK** on the Install a License dialog box to install the license.

Repeat the process until you have installed all of your licenses.

Certificate Trust Management

Cisco Security Manager downloads ASA images and IPS packages from Cisco.com over HTTPS, which uses certificates for establishing trust. Beginning with version 4.4, Security Manager has a certificate trust management feature. This feature helps you with improved handling of Cisco.com certificates for both types of downloads:

- ASA image downloads. For detailed documentation on certificate trust management for ASA image downloads, refer to Image Manager Page, page 11-41 or navigate to Tools > Security Manager Administration... > Image Manager > Help.
- IPS package downloads. For detailed documentation on certificate trust management for IPS package downloads, refer to Edit Update Server Settings Dialog Box, page 11-52 or navigate to Tools > Security Manager Administration... > IPS Updates > [Update Server group] > Edit Settings > Help.

Certificate Trust Management Feature

The certificate trust management feature in Security Manager has these characteristics:

- It behaves like a browser. It imparts trust to what you, as the user, consciously trust.
- It allows you to view the certificate and use your discretion in accepting it.
- It proactively validates a certificate to help you judge whether to accept or reject it. For example, it checks to see if a certificate is self-signed (not issued by a trusted Certificate Authority) and to see if it is expired, not yet valid, or revoked.
- After you accept a certificate, it stores that certificate on your Security Manager server.
- It provides transparency and control: You can retrieve and add a certificate, view a certificate, and remove a stored certificate.
- During communication with Cisco.com, it compares the live server certificate with the stored certificate and proceeds only upon a complete match. The complete certificate chain, not just the root certificate, is compared for a match. If there is a mismatch, the current operation is aborted until you view and accept the new certificate.
- It performs daily checks of your Security Manager server for certificate revocation and validity, and it removes any revoked or invalid certificates from your server. It does this by live contact with the CRL distribution points/URL present in the certificate. The default fixed schedule is for this daily check to be performed at 2:00 a.m.

Download Requirements

To download images from Cisco.com, you must retrieve, view, and accept both the latest image meta-data locator certificate and the latest certificate URL of the download site. The Security Manager interface has messages to assist you in key locations, and detailed documentation is available by referring to Image Manager Page, page 11-41 and Edit Update Server Settings Dialog Box, page 11-52.

Troubleshooting

During daily checks for certificate revocation and validity, the CRL revocation list is not stored on your Security Manager server. For that reason, if connectivity is lost, the daily checks fail to detect any possible certificate revocations. This problem will be solved after connectivity is restored.

If failure occurs while downloading ASA images or checking for IPS update packages, the most probable causes are the following:

- Site's certificate is not found on your Security Manager server

- There is a mismatch between the certificate received from the site and the stored certificate
- The site's certificate has expired

In all of these three cases listed above, the operation is aborted, and a message gives the cause of the error and the URL of the failed site. To recover, navigate to the user interface of the certificate feature (Tools > Security Manager Administration... > Image Manager or Tools > Security Manager Administration... > IPS Updates > [Update Server group] > Edit Settings); then retrieve, view, and accept the new certificate from the site and re-try the download.

If failure occurs while performing a check for IPS updates, verify that you have accepted both the certificate of the Cisco.com site used to obtain the meta-data information for IPS packages and the certificate of the actual download site of the IPS packages. Cisco recommends that you always configure email for notification of the job execution status. Then you can view the recommended actions in the email for recovering from the error. Copy the failed download URL from the email message to retrieve the certificate.

Because certificates are stored, if you upgrade to Security Manager 4.4 from a previous version, all communication with Cisco.com will fail. To resolve this problem, you must retrieve the certificates from the image meta-data locator and the download site URL.

If the stored certificate table in the user interface does not show the addition of a particular certificate, check to see if the daily checks for certificate revocation and validity have removed it because of revocation or expiration. You can do this by looking for the Certificate Revocation Check Task in the tomcat log; that log will enable you to determine the exact reason for the removal of the stored certificate.

Working with Audit Reports

When state changes occur in Security Manager, an audit entry is created in the audit log, which you can view by selecting **Manage > Audit Report**. The following topics provide more detailed information about audit reports:

- Understanding Audit Reports, page 10-19
- Generating the Audit Report, page 10-20
- Purging Audit Log Entries, page 10-23

Understanding Audit Reports

When state changes occur in Security Manager, an audit entry is created in the audit log, which you can view by selecting **Manage > Audit Report**.

The state changes that generate an event and create an audit entry are:

- Changes to the runtime environment:
  - System changes, such as login attempts (successful or failed), logout, and scheduled backups.
  - Authorization issues, such as failed attempts and security breaches.
  - Map changes, such as saving, deleting, and changing background map views.
  - Administrative changes, such as changing workflow modes.
- Changes to the state of Security Manager objects:
  - Activity changes, such as creating, editing, submitting, or approving an activity.

  - Deployment changes, such as creating, editing, or submitting a deployment job.
- Changes to the state of managed devices:
  - Object changes, such as changes to policy objects.
  - Inventory changes, such as adding, deleting, or modifying devices in the inventory.
  - Policy changes, such as creating, restoring, modifying, or deleting policies.
  - VPN changes, such as creating, modifying, or deleting a VPN.

When viewing the audit report, you can view subsets of entries by specifying search criteria to select only the desired records.

Related Topics

- Generating the Audit Report, page 10-20
- Purging Audit Log Entries, page 10-23

Generating the Audit Report

You can view the audit log to analyze the events that have occurred in the Security Manager System. This information can help you track changes that users have made to devices or to identify other system events of interest. The Audit Report window provides extensive search criteria to help you view the specific audit log entries that interest you.

Tip You can also view the audit logs through CiscoWorks Common Services. From the Common Services Server Administration page, select **Server > Reports**, and select **Audit Log** from the table of contents. Click **Generate Report** and you are presented with a list of logs, one for each day. Click the link for the desired log to open it. These logs are stored in the Program Files/CSCOpx/MDC/Logs/audit/ directory. For information about logging into Common Services, see Logging In to the Cisco Security Management Suite Server, page 1-12.

Related Topics

- Understanding Audit Reports, page 10-19

Step 1 Select **Manage > Audit Report** to open the Audit Report window.

Step 2 To reduce the report to a specific set of records that relate to an area of interest, enter the appropriate search criteria in the left pane and click **Search**. For detailed information about the search fields, see Using the Audit Report Window, page 10-21.

The following examples describe sample search criteria:

- To find out when the device router1 was removed from Security Manager management—Select **Devices > Delete** from the **Search by action** selector. In the **Search by all or part of the object name** field, enter the display name of the device (router1).
- To find out if a failed login attempt occurred in the system—Select **System > Authorization > Login > Failed** from the **Search by action** selector.

Step 3 To view the contents of an entry in the report, double click the entry. This action opens a dialog box where you can read the message related to the entry. You can scroll through the report in this dialog box by using the up and down arrow buttons.

Using the Audit Report Window

Use the Audit Report window to view records of state changes in Security Manager.

The Audit Report page contains two panes. Use the left pane to define the parameters for generating the audit report. The right pane displays the audit report using one row for each audit entry or message. The content of the audit report depends on the parameters you defined in the left pane. Therefore, all columns listed in the table might not be displayed in the generated audit report.

Navigation Path

Select **Manage > Audit Report**.

Related Topics

- Understanding Audit Reports, page 10-19
- Generating the Audit Report, page 10-20

Field Reference

Table 10-1 Audit Report Window

Element	Description
Search Criteria (Left Pane)	
The left side of the Audit Report window contains the search criteria for the report. The default report lists all state changes from yesterday and today, sorted with the most recent changes at the top.	
Search by action	One or more sources of actions to include in the audit report. If you do not make a selection, the report is not filtered based on action. You can select All to include all action sources.
Search by date	The time period to include in the report. Actions that occur between the from and to dates are displayed. Click the calendar icon to select the dates. This filter's default (reset position) is to include actions from yesterday to today.
Search for activity by state	This field works differently from the other search fields, and is primarily of use in Workflow mode. You can use this field to select one or more activities to include in the report. The activities are listed in the display box below the drop-down list. The drop-down list helps you find the activities on which you want to report. To use this search mechanism, select the activity state of the activities on which you want to report, and then select the activities. Use Ctrl+click to select multiple activities. Select No Activity to not filter by activity.

Table 10-1 Audit Report Window (continued)

Element	Description
Search by message warning level	The message warning level. The report is limited to messages of the selected severity. Use Ctrl+click to select multiple levels.
Search by user name	The username of the person who performed the actions to include in the report. To see Security Manager system-generated actions, enter the username System.
Search by a phrase in the message body	A string of text that should occur in the message of the audit report entries. You can enter a maximum of 1025 characters. The message is not visible in the report table. To see the messages related to an entry, double-click the entry.
Search by all or part of the object name	A string of text that should occur in the name of the object for which the audit entries were generated. You can enter a maximum of 1025 characters.
Search button	Click this button to generate the report in the right pane.
Reset button	Click this button to reset the search criteria and delete any values or selections you made.
Audit Report (Right Pane)	
The right side of the Audit Report window contains the audit report. Each row represents one audit entry. Double-click a row to open the Audit Message Details dialog box, where you can view a more readable layout of the information and to see the specific messages associated with the entry. You can scroll through the entries in the report from within the Audit Message Details dialog box.	
Message Level	The message warning level: Information, Warning, Success, Failure and Internal System Error.
Date	The date and time the action occurred.
Source	The origin of the audit entry: Objects, License, Admin, Firewall, Policy Manager, Devices, Topology, VPN, Config Archive, Deployment, System, and Activity.
Action	The action performed: Add, Assign, Create, Delete, Open, Purge, Unassign, and Update.
Object	The identifier of the object of the action. For example, if the category is device, then the object identifier could be the device name or IP address. If the category is deployment, then the object identifier could be job name, job ID, and so on. There frequently is no specific object name.
User Name	The username of the person performing the action.
Activity	The name of the activity in which the action occurred, if any.
# of rows per page	The number of rows to display on each page.
< arrow	Click this button to return to the previous page of the audit report.
> arrow	Click this button to advance to the next page of the audit report.

Purging Audit Log Entries

Security Manager automatically prunes the audit logs based on the age of the log entries. You do not need to actively manage the size of the log. However, you can change the defaults to increase or decrease the maximum size of the log and thus manage the overall size of the database.

To change the default settings for audit logs, select **Tools > Security Manager Administration** and select **Logs** from the table of contents (see Logs Page, page 11-62). The size of the log is controlled by the maximum number of days an entry can be, and the overall maximum number of entries that can be in the log. These settings work together, and entries are pruned on a periodic basis to keep the log to the maximum number of entries with none that are older than the maximum number of days. If you reduce the maximum size of the log, click **Purge Now** to delete the excess entries before the regular pruning cycle.

Note The Purge Now button only removes audit report entries from the database. It does not remove the *.csv files from the ***<*install_dir*>**\CSCOpx\MDC\log\audit folder. These *.csv files can be deleted directly.

You can also control the size of the log by changing the severity level of events that are captured in the log. For example, if you capture only Severe events, the log will probably remain small. However, reducing the level of information might reduce the value of the log.

Related Topics

- Understanding Audit Reports, page 10-19
- Generating the Audit Report, page 10-20
- Using the Audit Report Window, page 10-21

Taking Over Another User's Work

A user with administrative privileges can take over the work of another user in non-Workflow mode. Taking over another user's work is useful when a user is working on devices and policies, causing the devices and policies to be locked, and another user needs access to the same devices and policies.

Step 1 Select **Tools > Security Manager Administration** and select **Take Over User Session** from the table of contents to open the Take Over User Session page (see Take Over User Session Page, page 11-71).

Step 2 Select the user session you want to take over.

Step 3 Click **Take over session**. The changes made by the selected user are transferred to you. Any changes that have not already been committed are discarded.

If the selected user is logged in at the time changes are taken over, the user receives a warning message, loses the changes in progress, and then is logged out.

Changing Passwords for the Admin or Other Users

The admin user is a pre-defined user that has access to all Security Manager functions. When you install the product, you configure a password for the admin user. If you forget the password, you can use the following procedure to change it. You can also use this procedure to reset the password for other user accounts.

Step 1 Log into Windows on the Security Manager server and open a Windows command line window.

Step 2 Stop the daemon manager services by using this command:

net stop crmdmgtd

Tip You can also stop and start the daemon manager using the Services control panel.

Step 3 Run ResetPasswd.pl specifying admin as the user name. This example assumes you installed the product in the default directory; change the directory path if you used a different directory:

C:\Program Files\CSCOpx\bin\perl ResetPasswd.pl admin

You are prompted for a new password.

Tip If you want to change the password for a different user, replace **admin** with the desired user name.

Step 4 Start the daemon manager services by using this command:

net start crmdmgtd

Backing up and Restoring the Security Manager Database

You should regularly back up the Security Manager database in case you need to recover your work.

Tip The Security Manager database backup does not include the event data store used by the Event Manager service. If you want to back up event management data, see Archiving or Backing Up and Restoring the Event Data Store, page 69-36.

The following topics describe how to back up and recover the Security Manager database:

- Backing Up the Server Database, page 10-25
- Restoring the Server Database, page 10-27

Backing Up the Server Database

Security Manager uses CiscoWorks Common Services facilities to back up and restore its database. In the Security Manager client, select **Tools > Backup** to open the CiscoWorks Common Services Backup page for creating a backup schedule. You should regularly back up the database so that you can recover it if necessary.

After completing a backup, Security Manager compresses it. If you configure an e-mail address on the CiscoWorks Common Services backup page, you will get notifications of the completion of the backup and compression processes. If you have problems with file compression, or if you do not want to compress the backups, you can turn off backup compression. Edit the backup.properties file in the %NMSROOT%\conf folder (typically C:\Program Files\CSCOpx\conf) and change the backup compression property to specify NO instead of YES, as follows: VMS_FILEBACKUP_COMPRESS=NO.

Tip The backup includes the configuration and reporting databases, but it does not include the event storage areas. You can exclude the reporting database by changing the SKIP_RPT_DB_BACKUP property value to YES in the backup.properties file. Even if you specify YES, the backup will include reports generated by report schedules. For information on backing up the event data store, see Archiving or Backing Up and Restoring the Event Data Store, page 69-36.

While backing up and restoring data, both Common Services and Security Manager processes will be shut down and restarted. Because Security Manager can take several minutes to fully restart, users might be able to start their client before the restart is complete. If this happens, they might see the message “error loading page” in device policy windows.

We strongly recommend you take a backup of your current system before restoring an older backup.

You cannot restore a backup from an earlier version of Security Manager if that backup contains any pending data, which is data that has not been committed to the database. Before upgrading to a new version of Cisco Security Manager, we recommend that you commit or discard all uncommitted changes and then create a backup of your database. You can use the following instructions to help with committing or discarding pending data:

- **In non-Workflow mode:**
  - To commit changes, select **File > Submit**.
  - To discard uncommitted changes, select **File > Discard**.

  If there are multiple users with pending data, the changes for those users must also be committed or discarded. If you need to commit or discard changes for another user, you can take over that user’s session. To take over a session, select **Tools > Security Manager Administration > Take Over User Session**, select the session, and click **Take Over Session**.
- **In Workflow mode:**
  - To commit and approve changes, select **Manage > Activities**. From the Activity Manager window, select an activity and click **Approve**. If you are using an activity approver, click **Submit** and have the approver approve the activity.
  - To discard uncommitted changes, select **Manage > Activities**. From the Activity Manager window, select an activity and click **Discard**. Only an activity in the Edit or Edit Open state can be discarded.

You can also back up the database from a Windows command prompt using the following command:

[*path*]**perl** [*path*]**backup.pl** *backup_directory* [*log_filename* [**email=***email_address* [*number_of_generations* [**compress**]]]]

Syntax

[*path*]**perl** [*path*]**backup.pl**	The Perl script command. Include the path to the perl command and the backup.pl file if the path is not defined in the system path variable. The typical path for both is C:\Progra~1\CSCOpx\bin\.
backup_directory	The directory where you want to create the backup. For example, C:\Backups.
log_filename	(Optional.) The log file for messages generated during backup. Include the path if you want it created somewhere other than the current directory. For example, C:\BackupLogs. If you do not specify a name, the log is %NMSROOT%\log\dbbackup.log.
email=*email_address*	(Optional.) The email address where you want notifications sent. If you do not want to specify an email address, but you need to specify a subsequent parameter, enter **email** without the equal sign or address. You must configure SMTP settings in CiscoWorks Common Services to enable notifications. For more information, see Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.
number_of_generations	(Optional.) The maximum number of backup generations to keep in the backup directory. When the maximum is reached, old backups are deleted. The default is 0, which does not limit the number of generations kept.
compress	(Optional.) Whether you want the backup file to be compressed. If you do not enter this keyword, the backup is not compressed if VMS_FILEBACKUP_COMPRESS=NO is specified in the backup.properties file. Otherwise, the backup is still compressed. We recommend compressing backups.

Example

The following command assumes that you are in the directory containing the perl and backup.pl commands. It creates a compressed backup and log file in the backups directory and sends notifications to admin@domain.com. Note that you must specify a backup generation to include the compress parameter; if you specify any parameter after the log file parameter, you must include values for all preceding parameters.

```
perl backup.pl C:\backups C:\backups\backup.log email=admin@domain.com 0 compress
```

Tip If you stop a backup that is in progress, you need to delete the backup.LOCK file in the Security Manager installation directory (typically C:\Progra~1\CSCOpx) before you can perform another backup.

Restoring the Server Database

You can restore your database by running a script from the command line. You have to shut down and restart CiscoWorks while restoring data. This procedure describes how you can restore the backed up database on your server. A single backup and restore facility exists to back up and restore all applications installed on a CiscoWorks server; you cannot back up or restore individual applications.

If you install the applications on multiple servers, ensure that you recover the database backup that contains data appropriate for the installed applications.

Tip You can restore backups taken from previous releases of the application if the backup is from a version supported for direct local inline upgrade to this version of the application. For information on which versions are supported for upgrade, see the *Installation Guide for Cisco Security Manager* for this release of the product.

Step 1 Stop all processes by entering the following at the command line:

net stop crmdmgtd

Step 2 Restore the database by entering:

*$NMSROOT***bin\perl** *$NMSROOT***bin\restorebackup.pl** [**-t** *temporary_directory*] [**-gen** *generationNumber*] **-d** *backup_directory*[**-h**]

where:

- *$NMSROOT*—The full pathname of the Common Services installation directory (the default is C:\Program Files\CSCOpx).
- **-t** *temporary_directory*—(Optional) This is the directory or folder used by the restore program to store its temporary files. By default this directory is *$NMSROOT*\tempBackupData.
- **-gen** *generationNumber*—(Optional.) The backup generation number you want to recover. By default, it is the latest generation. If generations 1 through 5 exist, then 5 will be the latest.
- **-d** *backup_directory*—The backup directory that contains the backup to restore.
- **-h**—(Optional) Provides help. When used with -d BackupDirectory, help shows the correct syntax along with available suites and generations.

For example, to restore the most recent version from the c:\var\backup directory, enter the following command:

C:\Progra~1\CSCOpx\bin\perl C:\Progra~1\CSCOpx\bin\restorebackup.pl -d C:\var\backup

Tip If you are restoring a database that contains RME data, you might be asked if you want to collect inventory data. Collecting this data can take a long time. You might want to respond No and then configure RME to schedule an inventory. In RME, select **Devices > Inventory**.

Step 3 Examine the log file, *NMSROOT***log\restorebackup.log**, to verify that the database was restored.

Step 4 Restart the system by entering:

net start crmdmgtd

Step 5 If you restore a database that was backed up prior to installing a Security Manager service pack, you must reapply the service pack after restoring the database.

Generating Data for the Cisco Technical Assistance Center

Cisco Technical Assistance Center (TAC) personnel might ask you to submit a variety of data to help them identify and resolve problems you might encounter when using the application. The following topics can help you generate the required information. However, you should perform these tasks only at the direction of the TAC, because the information is not always required to resolve a problem.

- Creating Diagnostics Files for the Cisco Technical Assistance Center, page 10-28
- Generating Deployment or Discovery Status Reports, page 10-30
- Generating a Partial Database Backup for the Cisco Technical Assistance Center, page 10-30

Creating Diagnostics Files for the Cisco Technical Assistance Center

Cisco Technical Assistance Center (TAC) personnel may ask you to submit system configuration information in the form of a diagnostics file when you submit a problem report. The diagnostics file assists them with diagnosing the problem. You do not need to submit a diagnostics file unless they ask you to.

Before you create the diagnostics file, perform the actions that lead to the problem in your report. If necessary, you can control the level of detail in the diagnostics file by changing the settings on the Debug Options page (**Tools > Security Manager Administration > Debug Options**); see Debug Options Page, page 11-11.

Beginning with Version 4.7, Cisco Security Manager supports diagnostics in a new, light variant; this "Light Diagnostics" variant collects only basic information; as a result, the diagnostics file is smaller, and its generation is faster. The existing "General Diagnostics" variant is the same in 4.7 as it was in 4.6 and earlier versions.

General Diagnostics File

The general diagnostics file (CSMDiagnostics.zip) contains the following files and information:

- Configuration files
- Apache configuration and log files
- Tomcat configuration and log files
- Installation, audit, and operation log files
- The CiscoWorks Common Services Registry subtree ([HKEY_LOCAL_MACHINE][SOFTWARE][Cisco][MDC])
- Windows System Event and Application Event log files
- Host environment information (operating system version and installed service packs, amount of RAM, disk space on all volumes, computer name, and virtual memory size)

To create CSMDiagnostics.zip by using the GUI, follow this procedure:

1. Using the Security Manager client, select **Tools > Security Manager Diagnostics... > General Diagnostics...**. A dialog box is opened.

2. Click **OK** to start the file generation. The dialog box displays the progress.
3. When the file is generated, click **Close**.

To create CSMDiagnostics.zip by using the CLI, follow this procedure:

1. Open a command line window on the Security Manager server.
2. Run the **~MDC\bin\CSMDiagnostics** program.
3. CSMDiagnostics.zip is placed in the <*installation_location*>/MDC/etc folder, where <*installation_location*> is the drive and directory in which you installed CiscoWorks Common Services. The default value for <*installation_location*> is C:\Program Files (x86)\CSCOpx.
4. If you want to, you can specify a different folder where CSMDiagnostics.zip is placed. For example, you can specify **CSMDiagnostics C:\Temp**.
5. You should move or rename CSMDiagnostics.zip after you create it, because it will be overwritten the second time you generate it, with only one previous version (appended with "_old") being saved.

Note When creating CSMDiagnostics.zip by using the CLI, you must allow the command to complete before closing the window; if you do not, subsequent attempts to run **CSMDiagnostics** will not work properly. If you mistakenly close the window, delete the C:\Program Files\CSCOpx\MDC\etc\mdcsupporttemp folder before attempting to use the command again.

Light Diagnostics File

The light diagnostics file (CSMDiagnostics_light.zip) contains a subset of the general diagnostics file (CSMDiagnostics.zip), so it is smaller and generated faster.

To create CSMDiagnostics_light.zip by using the GUI, follow this procedure:

1. Using the Security Manager client, select **Tools > Security Manager Diagnostics... > Light Diagnostics...**. A dialog box is opened.
2. Click **OK** to start the file generation. The dialog box displays the progress.
3. When the file is generated, click **Close**.

To create CSMDiagnostics_light.zip by using the CLI, follow this procedure:

1. Open a command line window on the Security Manager Server.
2. Run the following command: ***<installation_location>***\MDC\diagnostics\script>rundiag.bat, where <*installation_location*> is the drive and directory in which you installed CiscoWorks Common Services. The default value for <*installation_location*> is C:\Program Files (x86)\CSCOpx.
3. Be certain that you run the command with the following 3 parameters and in the order shown:

   3.1. Installation Folder—Folder where Security Manager server is installed. This must not be changed or modified. An error in the path leads to failure in the generation of the diagnostics file. Example: C:\PROGRA~2\CSCOpx\MDC

   3.2. Destination Folder—Folder where the diagnostics file will be placed after generation. You can specify any path and folder where you want the file to be saved. If you want the file to be saved in the default path, then you must specify the default path explicitly. If you do not specify the path, there will be an error in generation. Example: C:\PROGRA~2\Light_Diagnostics

   3.3. "LightDiagnostics" string without a space. The string alphabets are not case-sensitive, so you can use capital or lower-case letters, but you must not use a space in this string. If you neglect to specify this string, part of the general diagnostics (i.e., not light diagnostics) logs will automatically be collected in the destination folder.

4. Complete command screen example:

   C:\Program Files (x86)\CSCOpx\MDC\diagnostics\script>rundiag.bat
   C:\PROGRA~2\CSCOpx\MDC C:\PROGRA~2\Light_Diagnostics LightDiagnostics

Generating Deployment or Discovery Status Reports

You can generate status reports for deployment and policy discovery jobs. If you have problems involving deployment or discovery, these reports can help the Cisco Technical Support (TAC) personnel resolve your problem. Although the reports are mainly intended for troubleshooting, you can also generate these reports for your own uses.

The status reports are generated as Adobe Acrobat (PDF) files on your workstation (you are prompted to select the location to save the PDF file). The report includes a summary of the job and summaries for each device in the job. Deployment status reports also include the full and delta configurations and the transcript of communications between Security Manager and the device.

You can generate deployment or discovery reports in the following ways:

- **Deployment Status Reports**
  - When a deployment job finishes, either successfully or unsuccessfully, by clicking the **Generate Report** button on the Deployment Status dialog box. See Deployment Status Details Dialog Box, page 8-32.
  - For previously completed jobs, by selecting the job in the Deployment Manager and clicking the **Generate Report** button. See Deployment Manager Window, page 8-16.
- **Discovery Status Reports**
  - During a discovery job, which might occur when adding devices or when rediscovering policies on devices already in the inventory, by clicking the **Generate Report** button on the Discovery Status dialog box. See Discovery Status Dialog Box, page 5-23.
  - For previously completed jobs, by selecting the job in the Policy Discovery Status dialog box and clicking the **Generate Report** button. See Policy Discovery Status Page, page 5-25.

Generating a Partial Database Backup for the Cisco Technical Assistance Center

Caution This topic explains how to create a partial database backup. Partial backups are incomplete and you cannot use them as a replacement for full backups. Partial backups are strictly for use in troubleshooting and you should generate one only if instructed to do so by Cisco Technical Assistance (TAC) personnel.

Partial database backups have the same characteristics as regular backups, but they contain a more limited set of data. When creating a partial backup, you are asked whether you want to include data from the Configuration Archive, and if you say yes, how many archive versions (per device) you want to include (regular backups include the entire Configuration Archive). For a description of regular backups, see Backing Up the Server Database, page 10-25.

Tip The partial backup will include or exclude the reporting database based on the settings in the backup.properties file as described in Backing Up the Server Database, page 10-25.

While backing up and restoring data, both Common Services and Security Manager processes will be shut down and restarted. Because Security Manager can take several minutes to fully restart, users might be able to start their client before the restart is complete. If this happens, they might see the message "error loading page" in device policy windows. Note that if you try to restore a partial backup, the system will point out that it is a partial backup, and you will have to confirm that you want to restore a partial backup.

To generate a partial backup, use the following command at a Windows command prompt on the Security Manager server.

[*path*]**perl** [*path*]**partial_backup.pl** *backup_directory* [*log_filename* [**email=***email_address* [*number_of_generations* [**compress**]]]]

Syntax

[*path*]**perl** [*path*]**partial_backup.pl**	The Perl script command. Include the path to the perl command and the partial_backup.pl file if the path is not defined in the system path variable. The typical path for both is C:\Progra~1\CSCOpx\bin\.
backup_directory	The directory where you want to create the backup. For example, C:\Backups.
log_filename	(Optional.) The log file for messages generated during backup. Include the path if you want it created somewhere other than the current directory. For example, C:\BackupLogs. If you do not specify a name, the log is %NMSROOT%\log\dbbackup.log.
email=*email_address*	(Optional.) The email address where you want notifications sent. If you do not want to specify an email address, but you need to specify a subsequent parameter, enter **email** without the equal sign or address. You must configure SMTP settings in CiscoWorks Common Services to enable notifications. For more information, see Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.
number_of_generations	(Optional.) The maximum number of backup generations to keep in the backup directory. When the maximum is reached, old backups are deleted. The default is 0, which does not limit the number of generations kept.
compress	(Optional.) Whether you want the backup file to be compressed. If you do not enter this keyword, the backup is not compressed if VMS_FILEBACKUP_COMPRESS=NO is specified in the backup.properties file. Otherwise, the backup is still compressed. We recommend compressing backups.

Example

The following command assumes that you are in the directory containing the perl and partial_backup.pl commands. It creates a compressed partial backup and log file in the backups directory and sends notifications to admin@domain.com. Note that you must specify a backup generation to include the compress parameter; if you specify any parameter after the log file parameter, you must include values for all preceding parameters. Also note that you are asked whether you want to include the Configuration Archive, and if yes, how many archive versions to include in the backup. In this example, five archive versions per device are included in the backup.

```
perl partial_backup.pl C:\backups C:\backups\pbackup.log email=admin@domain.com 0 compress
```

```
Root: c:\backups
Do you also want to take config-archive backup(Yes/No): Yes
How many previous config-archive you want to restore: 5
```

Configuring Security Manager Administrative Settings

Security Manager has default settings for many system functions that you can change if they do not fit the needs of your organization. To view and change these settings, select **Tools > Security Manager Administration**. You can then select items from the table of contents on the left of the window to view the default settings related to that item.

On most pages, when you change a setting, you must click **Save** to save your changes. If you make a mistake, you can click **Reset** to return the values to the previously saved values. You can also click **Restore Defaults** to return the settings to the Security Manager defaults.

Besides the pages that contain system defaults, the Security Manager Administration window includes items that relate to system administration activities, such as taking over another user's work or obtaining access to pages in Common Services to perform server security tasks.

The following topics describe the settings and actions available on each of the pages available in the Security Manager Administration window:

- API Settings Page, page 11-2
- AutoLink Settings Page, page 11-3
- CCO Settings Page, page 11-4
- Configuration Archive Page, page 11-6
- CS-MARS Page, page 11-7
- CSM Mobile Page, page 11-9
- Customize Desktop Page, page 11-10
- Debug Options Page, page 11-11
- Deployment Page, page 11-13
- Device Communication Page, page 11-21
- Device Groups Page, page 11-24
- Discovery Page, page 11-25
- Event Management Page, page 11-27
- Health and Performance Monitor Page, page 11-36
- Report Manager Page, page 11-38
- Identity Settings Page, page 11-38
- Image Manager Page, page 11-41

- IP Intelligence Settings Page, page 11-41
- Eventing Notification Settings Page, page 11-45
- IPS Updates Page, page 11-47
- ISE Settings Page, page 11-56
- Licensing Page, page 11-57
- Logs Page, page 11-62
- Policy Management Page, page 11-64
- Policy Objects Page, page 11-66
- Process Monitoring Settings Page, page 11-67
- Single Sign-on Configuration Page, page 11-68
- Rule Expiration Page, page 11-69
- Server Security Page, page 11-70
- Take Over User Session Page, page 11-71
- Ticket Management Page, page 11-72
- Token Management Page, page 11-73
- VPN Policy Defaults Page, page 11-74
- Workflow Page, page 11-75
- Wall Settings Page, page 11-77

API Settings Page

The Security Manager API settings page enables you to enable or disable the API service and change its settings.

Navigation Path

Click **Tools > Security Manager Administration** and select **API** from the table of contents.

Field Reference

Table 11-1 API Settings Page

Element	Description
Enable API Service	Whether to enable or disable the API service.
Result Set Page Size	Allowed values are 100 through 1000, inclusive.
Active client sessions	Allowed values are 1 through 10, inclusive.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

AutoLink Settings Page

The Security Manager Map view provides a graphical view of your VPN and layer 3 network topology. Using device nodes to represent managed devices and map objects to represent unmanaged objects such as devices, clouds, and networks, you can create topology maps with which to study your network. AutoLink settings enable you to exclude any one of five private or reserved networks from Map view. For example, you might want to exclude any test networks that are not relevant to the management tasks you are using Security Manager to perform.

Navigation Path

Click **Tools > Security Manager Administration** and select **AutoLink** from the table of contents.

Related Topics

- Creating and Managing Layer 3 Links on the Map, page 35-19
- Displaying Your Network on the Map, page 35-14

Field Reference

Table 11-2 AutoLink Page

Element	Description
Enable AutoLink for 10.0.0.0/8 Enable AutoLink for 172.16.0.0/12 Enable AutoLink for 192.168.0.0/16	Whether to automatically include or omit (deselected) these private networks from the maps you create.
Enable AutoLink for 127.0.0.0/8	Whether to automatically include or omit (deselected) the loopback network from the maps you create.
Enable AutoLink for 224.0.0.0/4	Whether to automatically include or omit (deselected) the multicast networks from the maps you create.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

ACL Hit Count Settings Page

The Security Manager ACL Hit Count Settings page enables you to configure and change the settings for Hit Count. This feature is available in Security Manager version 4.9 and higher for ASA and ASASM devices.

Navigation Path

Click **Tools > Security Manager Administration** and select **ACL Hit Count Settings** from the table of contents.

Field Reference

Table 11-3 Hit Count Settings

Element	Description
Hit Count History Persist Limit (per ACE)	The Hit Count History Persist Limit is the limit for the Hit Count history details that can be stored for the particular ACE in the database. The default value is 5 and you can enter a maximum value of 10.
Purge Scheduler Process Time	The Purge Scheduler Process Time is used by the Hit Count Scheduler to schedule the Hit Count purge job at the given time on a daily basis. Select a time from the drop-down list. The default is 12 AM.

CCO Settings Page

Use the CCO Settings page to configure the settings used to connect to Cisco.com.

Also, use the CCO Settings page for certificate trust management. (Security Manager downloads ASA images from Cisco.com over HTTPS, which uses certificates for establishing trust.) The certificate trust management feature on the Image Manager page is new in Security Manager 4.4. It will help you with improved handling of Cisco.com certificates for ASA image downloads:

- You can use it to view a certificate and use discretion in accepting it.
- After you accept a certificate, it is stored on your Security Manager server.
- You can see all your certificates in a summary table on the Image Manager page, and you can use that table to view or remove certificates.

Tip Please be sure to refer to "Retrieve Certificate" in the table below.

For detailed documentation of the certificate trust management feature, refer to Certificate Trust Management, page 10-18.

Navigation Path

Select **Tools > Security Manager Administration** and select **CCO Settings** from the table of contents.

Field Reference

Table 11-4 CCO Settings Page

Element	Description
Use IPS Updates Settings	If checked, the other settings on this page are disabled and the default prevails (the Cisco.com credentials from the IPS Updates page apply). **Caution** If checked, be sure that the Cisco.com credentials from the IPS Updates page are configured correctly. On that page, the default value for "Update From:" is "Local Server." You must choose "Cisco.com" to see certificate settings. Improper or incomplete certificate setting will prevent connectivity to Cisco.com, and all Cisco.com-related operations in this area will fail.
Username	The username Security Manager should use to log in to Cisco.com.
Password Confirm	The password for the username. Enter the password in both fields.
Proxy Server Settings	
Enable Proxy	Enables Image Manager to connect to Cisco.com via a web proxy server. When Enable Proxy is selected, the other proxy fields (including IP or Hostname, Port, Username, and Password) are enabled and used to connect to the web proxy.
Test Connection	Used to test for connectivity and credentials for Cisco.com.
Certificate	
Contact URL	• When selected, "Image Meta-data Locator" is used. This is the URL on Cisco.com that is used to obtain meta-data information about images. Meta-data information consists of the images applicable to a particular product, name, size, checksum, and URL to download for each image. • When selected, "Other" is used. You can enter any valid HTTPS URL. This URL is intended primarily for the HTTPS URL to download the image as obtained from the meta-data information about the image. This URL may be different from the URL of the image meta-data locator described in the previous paragraph; the certificate may be different, as well. **Caution** If you choose "Other," you need to explicitly add "https://dl.cisco.com" [without the quotation marks]: Enter it in the text field adjacent to the "Other" button. Failure to do this will prevent connectivity to Cisco.com, and all Cisco.com-related operations in this area will fail.

Table 11-4 CCO Settings Page (continued)

Element	Description
Retrieve Certificate	Used to connect to and retrieve the certificate from the selected ‘"Contact URL’." After retrieving the certificate it opens the Certificate Verification dialog, which along with a brief summary of the certificate, i.e., who the certificate is issued to, by whom, and the validity period of the certificate, gives you the following choices: • **View Certificate**—Opens the Certificate Viewer, where you can see all the details of the certificate: Certificate Authority, version, serial number, thumbprint, and other details. It shows the complete certificate chain information all the way up to the root issuing certificate Authority. • **Accept**—Accepts the certificate and adds it to the Cisco Security Manager. • **Reject**—Rejects the certificate and no action is taken. • **Cancel**—Closes the Certificate Verification dialog with no action taken. You must view and accept the following recommended certificates: • https://www.cisco.com • https://www.dl3.cisco.com • https://www.cloudsso.cisco.com • https://www.api.cisco.com
Certificates	A table that displays, for each certificate in your Security Manager installation, Subject, Issued By, and Accepted By.
View	Opens the Certificate Viewer for a certificate selected in the Certificate table.
Remove	Removes a certificate selected in the Certificate table.
Save button	Saves your changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Configuration Archive Page

Use the Configuration Archive page to define the default settings for the Configuration Archive tool, including how many configuration versions to save and the TFTP server to use for rolling back Cisco IOS software device configurations.

Navigation Path

Click **Tools > Security Manager Administration** and select **Configuration Archive** from the table of contents.

Related Topics

- Configuration Archive Window, page 8-23

- Rolling Back Configurations, page 8-63

Field Reference

Table 11-5 Configuration Archive Page

Element	Description
Max. Versions per Device Purge Now button Enable Configuration Archive Versions Auto Purge	The number of configuration versions you want to retain for each managed device, from 1 to 100. If you reduce the number, you can click **Purge Now** to immediately delete extra versions. Security Manager automatically deletes extra versions during its normal cleanup cycle if you select the **Enable Configuration Archive Versions Auto Purge** option.
TFTP Server for Rollback	The fully-qualified DNS hostname or IP address of the server to use for TFTP file transfers. TFTP is used during rollback for IOS devices when the configuration cannot be updated using the **configure replace** command, which does not force a system reload. Enter **localhost** to use the Security Manager server. By default, a TFTP server is enabled on the Security Manager server. If you specify a remote TFTP server, you must configure that server appropriately to provide TFTP services.
TFTP Root Directory	The root directory for configuration file transfers if you are using the Security Manager server as the TFTP server. Click **Browse** to select a directory on the Security Manager server. If you specify a server other than the Security Manager server as the TFTP host, Security Manager always uses the root directory of that TFTP server. You cannot specify a non-root directory for remote TFTP servers.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

CS-MARS Page

Use the CS-MARS page to register the Cisco Security Monitoring, Analysis and Response System servers that are monitoring your devices with Security Manager. By registering your CS-MARS servers, you can view messages and events captured in CS-MARS based on a device's firewall access rules or IPS signature rules configured in Security Manager. You must register a CS-MARS server before users can see events collected from it.

Tip If you are using CS-MARS global controllers, add them instead of the individual local controllers. By adding global controllers, Security Manager can identify the correct local controller for a device automatically, without you having to add each of the local controllers. This simplifies your CS-MARS configuration in Security Manager.

Navigation Path

Select **Tools > Security Manager Administration** and select **CS-MARS** from the table of contents.

Related Topics

- Registering CS-MARS Servers in Security Manager, page 72-38

Field Reference

Table 11-6 CS-MARS Page

Element	Description
CS-MARS Devices	The CS-MARS servers that are registered with Security Manager. • To add a server, click the Add (+) button and fill in the New or Edit CS-MARS Device Dialog Box, page 11-8. • To edit a server, select it and click the Edit (pencil) button. • To delete a server, select it and click the Delete (trash can) button. When you delete a server, the device properties for all devices that use the server are updated to remove the server connection. If a device is also monitored by another CS-MARS server on the list, its properties are updated to point to the other server.
When Launching CS-MARS Allow User to Save Credentials	The type of credentials Security Manager should use to log into CS-MARS when obtaining event information: • **Prompt users**—When the user tries to get event information from CS-MARS, prompt the user to log into CS-MARS. If you select this option, you can also select **Allow User to Save Credentials**, which gives users the option to save their credentials so they do not have to log into CS-MARS again the next time they request event status. • **Use CS-Manager Credentials**—When the user tries to get event information from CS-MARS, log into CS-MARS using the same username and password the user used to log into Security Manager.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.

New or Edit CS-MARS Device Dialog Box

Use the New or Edit CS-MARS Device dialog box to register a CS-MARS server with Security Manager. Users can obtain messages or event status for a device's firewall or IPS policies from the CS-MARS server that is monitoring the device. For more information, see Registering CS-MARS Servers in Security Manager, page 72-38.

Navigation Path

From the CS-MARS Page, page 11-7, click the Add button to add a new server, or select a server and click the Edit button.

Field Reference

Table 11-7 Add or Edit CS-MARS Device Dialog Box

Element	Description
CS-MARS Hostname/IP	The IP address or fully-qualified DNS host name of the CS-MARS server. **Tip** If you add a CS-MARS global controller, do not add any of the local controllers that the global controller monitors. Security Manager will automatically determine the local controller that is monitoring a specific device. Adding global controllers simplifies your CS-MARS configuration.
Username Password User Type	The username and password for logging into the server to validate that the CS-MARS server is running the appropriate software version and to obtain other basic information. Security Manager also uses this account to determine which CS-MARS server is monitoring a particular device. For CS-MARS local controllers, you can enter either a global or local user account. For global controllers, you must enter a global account. Identify the type of account in the User Type field.
Certificate Thumbprint Retrieve From Device button	The CS-MARS server certificate, a hexadecimal string that is unique to the device. Click **Retrieve From Device** to have Security Manager retrieve the certificate from the CS-MARS server. If the certificate is retrieved successfully, it is displayed. After verifying the certificate, click **Accept** to save it on the Security Manager server. You must have a correct certificate to use the CS-MARS server from Security Manager.

CSM Mobile Page

Use the CSM Mobile page of the Security Manager Administration window to enable or disable the CSM Mobile feature in Cisco Security Manager. If the CSM Mobile feature is enabled, users can access device health and summary information from mobile devices by navigating to the following link, where <SecManServer> is the DNS name or IP address of the Security Manager server:

https://<*SecManServer*>/mobile/

or

https://<*SecManServer*>/mobile

For more information about the types of information provided, see Dashboard Overview, page 72-1.

For more information about CSM Mobile, see CSM Mobile, page 72-11.

Navigation Path

Click **Tools > Security Manager Administration** and select **CSM Mobile** from the table of contents.

Field Reference

Table 11-8 CSM Mobile Page

Element	Description
Enable CSM Mobile Feature	Lets you enable or disable the CSM Mobile feature. If you disable this feature, you cannot access device health summary information from mobile devices.
Save button	Saves and applies changes. If you change whether the service is enabled, it stops or starts, as appropriate. You are shown a progress indicator.
Reset button	Resets changes to the last saved values.

Customize Desktop Page

Use the Customize Desktop page to control whether Security Manager applications close automatically after being idle for a specified time, to reset whether you are prompted to verify your actions in certain circumstances, and to control whether certain file operations can be performed on the Security Manager client.

Navigation Path

Select **Tools > Security Manager Administration** and select **Customize Desktop** from the table of contents.

Related Topics

- Installing Security Manager License Files, page 10-16
- Importing Policies or Devices, page 10-13
- Exporting the Device Inventory from the Security Manager Client, page 10-6
- Exporting Shared Policies, page 10-12
- Selecting IPS License Files, page 11-61

Field Reference

Table 11-9 Customize Desktop Page

Element	Description
Reset ‘Do Not Ask’ on Warnings button	Click this button to reestablish ‘Are you sure...?’ pop-up warnings. When you perform some actions, you are warned about the consequences and you are given the option to not be warned again. If you selected Do Not Ask Me Again for any of these warnings, clicking this button reenables the warning.

Table 11-9 Customize Desktop Page (continued)

Element	Description
Enable Idle Timeout Idle Timeout (minutes)	Whether to have the Security Manager client applications close automatically if you do not use them for the specified period of time. The timeout applies across all applications; working in one application resets the timer in all applications. If you select this option, enter the number of minutes that must elapse before closing the client in the Idle Timeout field. The default is to close the client after 120 minutes of inactivity.
Enable Client side file browser	Whether to allow file operations on the Security Manager client. If you select this option, you will be able to choose between client and server file systems when performing the following file operations: • Installing Security Manager license files • Installing IPS license files • Importing/exporting device inventory files • Importing/exporting shared policies • Creating the following file objects: – Cisco Secure Desktop Package – Plug-In—For browser plug-in files. – AnyConnect Profile – AnyConnect Image – Hostscan Image This option is enabled by default.
Global Search	
Enable Global Search	Whether to enable or disable the Global Search feature. This feature is enabled by default. **Tip** You can disable Global Search before performing bulk discovery or rediscovery of devices to improve performance. You can reenable Global Search and then recreate the index after discovery is complete or when users are least likely to be using the system.
Recreate Index	Click this button to regenerate the search index. The global search feature is not available while the index is being recreated.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Debug Options Page

Use the Debug Options page to configure the severity level of messages to include in debugging logs and to determine what other debugging information is collected.

You should change debugging levels only if the Cisco Technical Assistance Center (TAC) asks you to change them. This makes it possible for you to include more detailed information in the CSMDiagnostics.zip file.

After you change the message level for the appropriate subcomponent, redo the actions that are resulting in system problems. After the problems occur, create the CSMDiagnostics.zip file (or the CSMDiagnostics_light.zip file) by selecting **Tools > Security Manager Diagnostics... > General Diagnostics...** (or **Tools > Security Manager Diagnostics... > Light Diagnostics...**). You can then reset the debug options to the default levels so that the Security Manager server does not become bogged down collecting extra debug information. For more information about generating the CSMDiagnostics.zip file, see Creating Diagnostics Files for the Cisco Technical Assistance Center, page 10-28.

By default, logs contain messages of the Error severity or worse. The severity levels in order of severity are:

- Severe—Problems that make the system unusable.
- Error—Problems from which Security Manager cannot recover.
- Warning—Unexpected conditions from which Security Manager can recover.
- Info—Informational messages.
- Debug—Internal status information.

Navigation Path

Select **Tools > Security Manager Administration**, then select **Debug Options** from the table of contents.

Field Reference

Table 11-10 Debug Options Page

Element	Description
Capture Discovery/Deployment Debugging Snapshots to File	Whether Security Manager generates data files about configuration generation, deployment, and discovery as these functions are performed. The temporary data files are stored in the MDC\temp directory in the Security Manager installation folder on the server, and you can use these files for debugging purposes. Enable this setting if you encounter problems with deployment or discovery. **Note** Selecting this check box slows down Security Manager response time. Enable this option only in limited circumstances. If you need to send these files to Cisco TAC for debugging, encrypt the files, because they can contain sensitive data such as passwords. **Note** Do not delete any snapshot files under MDC\temp directory in the Security Manager installation folder while discovery (or) deployment is in progress. You can delete them when Security Manager is idle. In addition, ensure you are not deleting any of the default files.
Deployment Debug Level	The message severity level for deployment-related actions such as device communication.
Event Manager Debug Level	The message severity level for the Event Manager subsystem.

Table 11-10 Debug Options Page (continued)

Element	Description
Health and Performance Monitor Debug Level	The message severity level for the Health and Performance Monitor subsystem.
Image Manager Debug Level	The message severity level for the Image Manager subsystem.
Firewall Services Debug Level	The message severity level for firewall-related policies.
IOS Platform Debug Level	The message severity level for Cisco IOS Software platform policies.
PIX Platform Debug Level	The message severity level for PIX, ASA, and FWSM platform policies.
Report Manager Debug Level	The message severity level for the Report Manager subsystem.
VPN Services Debug Level	The message severity level for VPN services policies.
API Debug Level	The message severity level for the Application Programming Interface subsystem.
Save button	Saves your changes.
Reset button	Restores all fields to their previous values.
Restore Defaults button	Resets values to Security Manager defaults.

Deployment Page

Use the Deployment page to define the default methods by which Security Manager deploys configurations to devices. You can override some of these settings when you create deployment jobs.

Navigation Path

Select **Tools > Security Manager Administration** and select **Deployment** from the table of contents.

Related Topics

- Chapter 8, "Managing Deployment"
- Chapter 6, "Managing Policy Objects"

Field Reference

Table 11-11 Deployment Page

Element	Description
General Parameters	
Purge Debugging Files Older Than (days)	The maximum number of days the system should keep debugging files. Debug files are automatically deleted. If you decrease the number of days, you can click **Purge Now** to immediately delete all debugging files older than the number of days specified.

Table 11-11 Deployment Page (continued)

Element	Description
Default Deployment Method Directory	The method to use as the default method for deploying configurations to devices: • Device—Deploys the configuration directly to the device or to the transport mechanism specified for the device. For more information, see Deploying Directly to a Device, page 8-9. • File—Deploys the configuration file to a directory on the Security Manager server. If you select File, specify the directory to which you want to deploy the configuration file in the Destination column. Even if you select file as the default, the setting does not apply to IPS devices; you can use device deployment only for IPS devices. For more information, see Deploying to a File, page 8-11. You can override this method when you create deployment jobs.
When Out of Band Changes Detected	How Security Manager should respond when it detects that changes were made directly on the device CLI since a configuration was last deployed to the device. Out of band change detection works correctly only when deploying to device, not to file, and applies only when the deployment method is configured to obtain the reference configuration from the device (see below for a description of the Reference Configuration setting). This setting specifies the default action, which you can override when you create deployment jobs. You can choose one of the following: • **Overwrite changes and show warning** (default)—If changes were made to the device manually, Security Manager continues with the deployment, overwrites the changes, and displays a warning notifying you of this action. • **Cancel deployment**—If changes were made to the device manually, Security Manager cancels the deployment and displays a warning notifying you of this action. • **Do not check for changes**—Security Manager does not check for changes and deploys the changes to the device, overwriting any local modifications. For a more complete discussion of out-of-band change handling, see Understanding How Out-of-Band Changes are Handled, page 8-12. **Note** For devices in which failover is not configured, if you select the **Cancel Deployment** option when Out of Band changes are detected, the bootstrap configuration may cause deployments to fail. For deployments to be successful, you must configure failover before discovering the device in Security Manager.

Table 11-11 Deployment Page (continued)

Element	Description
Deploy to File Reference Configuration	The configuration that Security Manager uses to compare new policies against the previous configuration for the device, if you are deploying the configuration to a file on the Security Manager server. • Archive (default)—The most recently archived configuration. • Device—The current running device configuration, which is obtained from the device. After comparing the configurations, Security Manager generates the correct CLI for deployment.
Deploy to Device Reference Configuration	The configuration that Security Manager uses to compare new policies against the previous configuration for the device, if you are deploying the configuration directly to the device (or to a transport server). • Archive—The most recently archived configuration. • Device (default)—The current running device configuration, which is obtained from the device. After comparing the configurations, Security Manager generates the correct CLI for deployment.
Allow Download on Error	Whether deployments to devices should continue even if there are minor device configuration errors.
Save Changes Permanently on Device	Whether to save the running configuration to the startup configuration (using the **write memory** command) after deploying a configuration to a device. This applies to PIX, FWSM, ASA, or Cisco IOS devices. If you deselect this check box, the startup configuration is not changed, which means your configuration changes will be lost if the device reloads for any reason.
Preselect Devices with Undeployed Changes	Whether the list of changed devices you see when you create a deployment job has all changed devices preselected. If you deselect this option, users must manually select the devices to include in the deployment job.
Enable Auto Refresh in Deployment Main Panel	Whether the deployment job and schedule status information should be automatically refreshed in the Deployment Manager window. If you deselect this option, you must click the Refresh button to refresh the information manually.
Remove Unreferenced SSL VPN Files on Device (ASA Only)	Whether to have Security Manager delete files related to the SSL VPN configuration from the device if the files are no longer referred to by the device's SSL VPN configuration. If you deselect this option, unused files remain on the device after deployment.

Table 11-11 Deployment Page (continued)

Element	Description
Mask Passwords and Keys When Viewing Configs and Transcripts Mask Passwords and Keys When Deploying to File	The conditions, if any, under which Security Manager will mask the following items so that they cannot be read: passwords for users, enable mode, Telnet, and console; SNMP community strings; keys, including those for TACACS+, Preshared Key, RADIUS server, ISAKMP, failover, web VPN attributes, logging policy attributes, AAA, AUS, OSPF, RIP, NTP, logging FTP server, point-to-point protocol, Storage Key, single sign-on server, load balancing, HTTP/HTTPS proxy, and the IPSEC shared key. • Mask Passwords and Keys When Viewing Configs and Transcripts—This option affects only the screen display of the credentials, which guards against unauthorized personnel viewing them. If you do not select this option, credentials in full transcripts might still be masked depending on how the device handles them. • Mask Passwords and Keys When Deploying to File—This option affects the contents of configuration files that are deployed to file, making them undeployable to actual devices. Select this option only if you will never need to actually deploy these configurations to real devices. Selecting this option has no effect on whether credentials are masked when viewed.
Deploy only new or modified Flexconfigs	Whether to deploy FlexConfigs only one time after creation or modification of a FlexConfig, or to deploy all FlexConfigs with each deployment. This option is selected by default. **Note** If you have FlexConfigs that need to be deployed with each deployment, then you will need to disable this option. After changing this setting, you will need to manage one-time FlexConfigs by deleting them after they have been deployed.

Table 11-11 Deployment Page (continued)

Element	Description
ACL Parameters	
Optimize the Deployment of Access Rules For (IPv4 and IPv6 access rules.)	How firewall rules are deployed. You can choose one of the following: • Speed (default)—Increases deployment speed by sending only the delta (difference) between the new and old ACLs. This is the recommended option. By making use of ACL line numbers, this approach selectively adds, updates, or deletes ACEs at specific positions and avoids resending the entire ACL. Because the ACL being edited is still in use, there is a small chance that some traffic might be handled incorrectly between the time an ACE is removed and the time that it is added to a new position. The ACL line number feature is supported by most Cisco IOS, PIX and ASA versions, and became available in FWSM from FWSM 3.1(1). • Traffic—This approach switches ACLs seamlessly and avoids traffic interruption. However, deployment takes longer and uses more device memory before the temporary ACLs are deleted. First, a temporary copy is made of the ACL that is intended for deployment. This temporary ACL binds to the target interface. Then the old ACL is recreated with its original name but with the content of the new ACL. It also binds to the target interface. At this point, the temporary ACL is deleted. **Note** For FWSM devices, this option affects processing only if you also select the Let FWSM Decide When to Compile Access Lists option.
Firewall Access-List Names (IPv4 and IPv6 access rules.)	How ACL names are deployed to devices if the access rule does not have a name in Security Manager. • Reuse existing names—Reuse the ACL name that is configured in the reference configuration (which is usually from the device). • Reset to CS-Manager generated names—Reset the name to a Security Manager auto-generated ACL name.

Table 11-11 Deployment Page (continued)

Element	Description
Enable ACL Sharing for Firewall Rules (IPv4 and IPv6 access rules.)	Whether Security Manager should share a single access control list (ACL) for an access rule policy with more than one interface. If you do not select this option, Security Manager creates unique ACLs for every interface to which you apply an IPv4 or IPv6 access rule policy. The sharing of ACLs is done only for ACLs created by access rule policies. If you select this option, Security Manager evaluates the access rules policy for each interface and deploys the minimum number required to implement your policy while preserving your ACL naming requirements. For example, if you use an interface role to assign the same rules to four interfaces, you specify **Reset to CS-Manager generated names** for the Firewall Access-List Names property, and you do not specify ACL names for the interfaces in the access control settings policy, only a single ACL is deployed, and each interface uses that ACL. If you select this option, keep the following in mind: • An interface might use an ACL that is named for a different interface. • If you specify a name for the ACL in the access control settings policy, an ACL by that name is created even if it is otherwise identical to one used by another interface. Names specified in this policy have precedence over any other settings. • If you select **Reuse existing names** for the Firewall Access-List Names property, the existing names are preserved (unless you override them in the access control settings policy). This means that you might end up with duplicate ACLs under different names if duplicate ACLs already exist on the device. • Hit count statistics are based on ACL, not on interface, so a shared ACL provides statistics that are combined from all interfaces that share the ACL. • Sharing ACLs is primarily beneficial for memory-constrained devices such as the FWSM.
Let FWSM Decide When to Compile Access Lists (IPv4 access rules only.)	Whether to have the Firewall Services Module (FWSM) automatically determine when to compile access lists. Selecting this option might increase deployment speed but traffic might be disrupted and the system might become incapable of reporting ACL compilation error messages. If you select this option, you can use the Optimize the Deployment of Access Rules For Traffic setting to mitigate potential traffic disruptions. When deselected, Security Manager controls ACL compilation to avoid traffic interruption and to minimize peak memory usage on the device. ⚠ **Caution** You should not select this option unless you are experiencing deployment problems and you are an advanced user.

Table 11-11 Deployment Page (continued)

Element	Description
Remove Unreferenced Access-lists on Device (IPv4 and IPv6 access rules.)	Whether to delete any access lists that are not being used by other CLI commands managed by Security Manager from devices during deployment. **Note** After enabling this option, Security Manager will remove access lists during deployment that are not used in any policies managed or discovered by Security Manager. If any policy that is NOT discovered or managed by Security Manager is using such an access list, Security Manager will still attempt to delete that object during deployment. This also applies to access lists that are used in FlexConfigs but are not used in any other policies managed by Security Manager.
Generate ACL Remarks During Deployment (IPv4 and IPv6 access rules.)	Whether to display ACL warning messages and remarks during deployment.
Preserve Sections for Access Rules	Whether to deploy the section name under which access rules are organized. This option ensures that if a device is discovered or rediscovered, the section names will not be lost.
Generate CSM Rule Number	Whether to deploy the rule number used in the Cisco Security Manager user interface. This option helps in correlating an access rule in a device configuration to its position in rule table.
Object Group Parameters	
Remove Unreferenced Object Groups from Device (PIX, ASA, FWSM, IOS 12.4(20)T+) (IPv4 and IPv6 objects.)	Whether Security Manager should remove object groups that are not being used by other CLI commands managed by Security Manager from devices during deployment. Object groups include network/host, service, and identity user groups. **Note** After enabling this option, Security Manager will remove objects during deployment that are not used in any policies managed or discovered by Security Manager. If any policy that is NOT discovered or managed by Security Manager is using such an object, Security Manager will still attempt to delete that object during deployment. In such cases, deployment will fail with a transcript error indicating that it was unable to delete the object. **Tip** Network/host objects that include object NAT configurations on ASA 8.3+ devices are never considered unreferenced.

Table 11-11 Deployment Page (continued)

Element	Description
Create Object Groups for Policy Objects (PIX, ASA, FWSM, IOS 12.4(20)T+) Create Object Groups for Multiple Sources, Destinations or Services in a Rule (PIX, ASA, FWSM, IOS 12.4(20)T+) Optimize Network Object Groups During Deployment (PIX, ASA, FWSM, IOS 12.4(20)T+) (IPv4 and IPv6 objects.)	Whether Security Manager should create object groups, such as network objects, service group objects, and identity user group objects, to replace comma-separated values in a rule table cell for the indicated devices. When deselected, Security Manager flattens the object groups to display the IP addresses, sources and destinations, users, ports, and protocols for these devices. **Tip** These options do not apply to host, network, or address range network/host objects, or to service objects (as opposed to service group objects), which are always created as objects. Multiple FQDN network objects can be grouped into a single network object. If you select this option, you can also select these options: • Create Object Groups for Multiple Sources, Destinations or Services in a Rule—Whether to automatically create network objects, service objects, and identity user group objects to replace comma-separated values in a rule table cell that resulted when multiple rules were combined. The objects are created during deployment and are in the format of 'CSM_INLINE...' for example, 'CSM_INLINE_src_rule_8589960758'. For more information, see Combining Rules, page 12-22. • Optimize Network Object Groups During Deployment—Whether to optimize network object groups by making them more succinct. For more information on optimizing policy objects, see Optimizing Network Object Groups When Deploying Firewall Rules, page 12-35.
IPS Parameters	
Generate transcripts for IPS Auto-Update Jobs	
Attach transcripts to email for IPS Auto-Update Jobs	

Table 11-11 Deployment Page (continued)

Element	Description
Remove Unreferenced Signature and Event Action Variables from IPS Device (IPS Parameters object group)	Whether to delete the unused variables from the sensor (IPS device) configuration during the next deployment. IPS Event and Signature Variables are defined as policy objects in Security Manager. Disabled by default (checkbox is cleared by default); that is, do not remove the unreferenced variables. Applies to the following variables; applies to both IPv4 and IPv6: • signature source and destination addresses • signature service port variables in signature engine parameters • victim and attacker addresses in event action filters • network information target addresses Does not apply to the following variables: • signature source port • OS identification address • signature destination port
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Device Communication Page

Use the Device Communication page to define default settings for communicating with devices. These settings mainly affect device inventory and policy discovery and configuration deployment. You can override the transport settings for individual devices in the device properties for the device.

If you change the transport protocol settings, ensure that your devices are appropriately configured to accept those types of connections.

Navigation Path

Select **Tools > Security Manager Administration** and select **Device Communication** from the table of contents.

Related Topics

- Adding Devices to the Device Inventory, page 3-7
- Chapter 3, “Managing the Device Inventory”
- Chapter 2, “Preparing Devices for Management”
- Viewing or Changing Device Properties, page 3-41

Field Reference

Table 11-12 Device Communication Page

Element	Description
Device Connection Parameters	
Device Connection Timeout	The number of seconds that Security Manager has to establish a connection with a device before timing out.
Retry Count	The number of times that Security Manager should try to establish a connection to a device before concluding that the connection cannot be completed. The default value is 3.
Socket Read Timeout	For SSH and Telnet sessions, the maximum number of seconds Security Manager can wait for incoming data before concluding that the connection is lost.
Transport Protocol (IPS)	The default transport protocol for IPS sensors and routers that include the IPS feature. The default is HTTPS.
Transport Protocol (IOS Routers 12.3 and above)	The default transport protocol for routers that run Cisco IOS software release 12.3 and above. The default is HTTPS.
Transport Protocol (Catalyst Switch/7600)	The default transport protocol for Catalyst 6500/7600 devices and all other Catalyst switches, regardless of the Cisco IOS software version running on the devices. The default is SSH.
Transport Protocol (IOS Routers 12.2, 12.1)	The default transport protocol for routers that run Cisco IOS software releases 12.1 and 12.2. The default is Telnet.
Connect to Device Using	The type of credentials Security Manager should use when accessing devices. For more information, see Understanding Device Credentials, page 3-4. • **Security Manager User Login Credentials**—Security Manager contacts the device using the credentials that you entered while logging in to Security Manager. The same set of credentials are used for all devices regardless of the credentials configured for each device on the Device Credentials page. • **Security Manager Device Credentials**—Security Manager contacts the device using the credentials specified in the Device Properties Credentials page. This is the default. **Caution** You must use Security Manager Device Credentials, not Security Manager User Login Credentials, if a connection to an IPS sensor is involved. When Security Manager contacts an IPS sensor, it must use device credentials whether or not someone is logged in to Security Manager.
SSL Certificate Parameters	

Table 11-12 Device Communication Page (continued)

Element	Description
Device Authentication Certificates (IPS) Device Authentication Certificates (Router) PIX/ ASA/ FWSM Device Authentication Certificates Add Certificate button	How to handle device authentication certificates for SSL (HTTPS) communications. You can configure different behaviors for different types of devices, but the settings have the same meaning: • **Retrieve while adding devices**—Security Manager automatically obtains certificates from the devices while you add devices from the network or from an export file. • **Manually add certificates**—Security Manager does not automatically accept certificates from the device. Click **Add Certificate** to open the Add Certificate dialog box (see Add Certificate Dialog Box, page 11-24) where you can manually add the thumbprint before you try to add the device from the network or from an export file. You can also add certificates for devices that you create manually from the Device Properties Credentials page to be successful. For more information, see Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5. • **Do not use certificate authentication**—Security Manager ignores device authentication certificates. This option leaves your system vulnerable to third-party interference with device validation. We recommend that you do not use this option.
Accept Device SSL Certificate after Rollback	For devices that use SSL, whether to obtain the certificate installed on an IPS device, firewall device, FWSM, ASA, or Cisco IOS router from the device when you roll back the configuration on the device.
HTTPS Port Number	The default port number that the device uses for secure communication with Security Manager (as well as other management applications that use these protocols). This value overrides the HTTPS port number that you configure in the HTTP policy for a device. **Note** If you configure the local HTTP policy to be a shared policy and assign the HTTP policy to multiple devices, the HTTPS port number setting in the shared policy overrides the port number configured in the Device Properties Credentials page for all devices to which the policy is assigned. In addition to providing access to the device through the Cisco web browser user interface, the HTTPS port number is used by device management applications (such as the Cisco Router and Security Device Manager (SDM)) and monitoring tools to communicate with the device. **Note** The security appliance can support both SSL VPN connections and HTTPS connections for device manager administrative sessions simultaneously on the same interface. Both HTTPS and SSL VPN use port 443 by default. Therefore, to enable both HTTPS and SSL VPN on the same interface, you must specify a different port number for either HTTPS or WebVPN. An alternative is to configure SSL VPN and HTTPS on different interfaces.

Table 11-12 Device Communication Page (continued)

Element	Description
Overwrite SSH Keys	Whether Security Manager can overwrite the SSH key for a device when it changes on the device. For SSH connections, a correct key is required for successful communication. Deselect this check box with caution, and only if you require a greater level of security. Security Manager does not communicate with the device if keys are changed on the device.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Add Certificate Dialog Box

Use the Add Certificate dialog box to add device certificates manually for devices that use the SSL transport protocol (firewall devices, FWSMs, ASAs, IPS devices, and Cisco IOS devices). Adding the device certificates manually gives you the highest level of security because then an intruder is prevented from introducing a fraudulent certificate thumbprint. Device certificates are stored in the database to be used for device authentication.

For more information about manually adding SSL certificates, see Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5.

Navigation Path

Select **Tools > Security Manager Administration**, select **Device Communication** from the table of content, and click **Add Certificate**.

Field Reference

Table 11-13 Add Certificate Dialog Box

Element	Description
Host Name or IP Address	The hostname or IP address of the device for which you are adding the certificate.
Certificate Thumbprint	The certificate thumbprint, which is a string of hexadecimal digits that is unique to the device.

Device Groups Page

Use the Device Groups page to manage the device groups and group types defined in the device inventory.

Navigation Path

Select **Tools > Security Manager Administration**, then select **Device Groups** from the table of contents.

Related Topics

- Understanding Device Grouping, page 3-60
- Working with Device Groups, page 3-59

Field Reference

Table 11-14 Device Groups Page

Element	Description
Groups	Displays the device groups and group types. To rename a group or type, select it and then click it again to make the text editable. Type in the new name and press Enter.
Add Type button	Click this button to create a new group type. The type is added with a default name. Overtype the name and press Enter.
Add Group to Type button	Click this button to add a device group to the selected device group or group type.
Delete button (trash can)	Click this button to delete the selected device group or group type and all device groups that it contains. Deleting a device group or group type does not delete any devices it contains.
Save button	Saves your changes.
Reset button	Restores all fields to their previous values.

Discovery Page

Use the Discovery page to define how Security Manager should handle certain types of objects or events during inventory and policy discovery. You can also control how long Security Manager keeps discovery tasks.

Navigation Path

Select **Tools > Security Manager Administration** and select **Discovery** from the table of contents.

Field Reference

Table 11-15 Discovery Page

Element	Description
Prepend Device Name when Generating Security Context Names	Whether the name of the device that contains the security context should be added to the front of the security context's name. For example, if a security context is named admin, and it is contained in the device with the display name 10.100.15.16, the name that will appear in the Device selector is 10.100.15.16_admin. If you do not prepend the device name, the security context name appears in the inventory by itself. Because Security Manager does not place security contexts in a folder related to the parent device, the only way to easily see contexts that are related to a device is to prepend the device name. If you do not prepend device names, Security Manager adds a numbered suffix to distinguish identically named devices. For example, if the admin context exists in more than one firewall, you will see admin_01, admin_02, and so on, in the Device selector.
Purge Discovery Tasks Older Than	The number of days to save discovery and device-import tasks. Tasks older than the number of days you enter are deleted.
Reuse Policy Objects for Inline Values	Whether to substitute any named policy objects, such as network/host or identity user group objects already defined in Security Manager, for inline values in the CLI. For more information on policy objects, see Chapter 6, "Managing Policy Objects". **Tip** Although this option generally applies to network/host objects, it does not apply to FQDN network/host objects because you cannot specify a fully-qualified domain name (FQDN) as an inline value.
Allow Device Override for Discovered Policy Objects	For the types of objects for which overrides are possible, whether to allow users to override the parent object values at the device level for policy objects that are discovered. For example, if you select this option, if you run policy discovery on a device that has an ACL with the same name as an ACL policy object in Security Manager, the name of the discovered policy object is reused, but a device-level override is created for the object. If you deselect this option, a new policy object is created with a number appended to the name. **Tip** For objects that have subtypes, such as network/host and service, overrides are limited to within a type. For example, an override can be created for a network/host group when discovering a same-named network/host group, but no override would be created when discovering a same-named network/host address range. Instead, the newly-discovered object will have a number appended to the name. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18.

Table 11-15 Discovery Page (continued)

Element	Description
On Error, Rollback Discovery for Entire Device	Whether Security Manager should roll back all discovered policies if even one error is encountered for a single policy during policy discovery. When deselected, Security Manager keeps the policies successfully discovered and discards only those policies with errors. For more information on policy discovery, see Discovering Policies, page 5-12.
Auto-Expand Object Groups with Prefixes	Expands object groups, such as network or identity user group, with the listed prefixes during the device import process. Separate the prefixes with a comma. This expansion causes the elements of the object group to display as separate items in the discovered policies. For more information, see Expanding Object Groups During Discovery, page 12-35. **Tip** This option does not apply to policy objects created from the **object network** or **object service** commands from ASA 8.3+ devices. These commands create host, FQDN, network, or address range network/host objects or service objects.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Event Management Page

Use the Event Management page to enable event management, which allows you to view ASA, FWSM, and IPS events using the Event Viewer. You can also configure settings required for event collection.

The Event Manager service is also required by the Report Manager application, which allows you to view reports that aggregate information collected by the service.

Tip If you get a message that Event Viewer is unavailable when you select Launch > Event Viewer, but the **Enable Event Management** option is selected on this page, try restarting the Event Manager Service. First, deselect the Enable option and click Save. Wait for the service to stop. Then, select the Enable option, click Save, and wait for the service to finish restarting. You can then try opening Event Viewer again.

Navigation Path

Click **Tools > Security Manager Administration** and select **Event Management** from the table of contents.

Field Reference

Table 11-16 Event Management Page

Element	Description
Event Management Options	
Enable Event Management	Whether to enable the Event Manager service, which allows Security Manager to collect event information. If you disable this feature, you cannot use the Event Viewer or Report Manager applications. **Tip** If you change this setting and click **Save**, you are prompted to confirm that you want to start or stop the Event Manager Service. If you click Yes, the service is started or stopped immediately, and you are shown a progress indicator and told when the change is completed. Wait until the status change is completed before continuing.
Event Data Store Location	The directory to use for collecting event information. This is known as the primary event store. Click **Browse** to select a directory on the Security Manager server. If the directory does not yet exist, create it in Windows Explorer. You cannot create the directory from within Security Manager. **Tip** If you change the location after you have started using the Event Manager service, you cannot query old events.
Event Data Store Disk Size	The amount of disk space you want to allocate for storing event data, in gigabytes (GB). Events are incrementally deleted (rotated out) from the extended store when it becomes 90% full. Before changing this setting, consider the following: • If you reduce the size below the amount of disk space already used by event data, the oldest events are deleted until your new size limit is reached. • You can see a visual representation of the amount of space currently used for event data. Open the Event Viewer (**Launch > Event Viewer**), then from Event Viewer, select **Views > Show Event Store Disk Usage**.
Event Syslog Capture Port	The port on which you want to enable syslog event capture. The default is 514. You must ensure that the Security Manager server, and intervening firewalls, allow incoming traffic on this port for Security Manager to collect the events. Managed devices must be configured to send syslog information to this port on the Security Manager server. **Tip** If you change this port, you must also change the Syslog Servers policies for all ASA and FWSM devices and security contexts that send events to Security Manager. For more information, see Syslog Servers Page, page 54-27.
Event Data Pagination Size	The maximum number of events per page each query response can contain. The default is 20000, but you can select a different size from the list of supported values. **Note** In Security Manager 4.10, the maximum number of events per page has been increased to 100000.

Table 11-16 Event Management Page (continued)

Element	Description
Extended Store Management Options	
Auto Copy Events to Extended Store	Whether you want to define an extended storage location for event storage. Events are copied from the regular event storage location to the extended location so that they remain available for use. When you query for historical events in Event Viewer, events in the extended storage location are automatically retrieved if they are needed. **Tip** You are prompted to verify that you want to start the extended service and to make changes to the extended storage location.
Extended Data Store Location	The location of the extended data store for events. This location can be on directly-attached storage that appears as a drive on the server and that uses DAS protocols. For example, SAN storage attached through fiber channel. CIFS storage is not supported. Click **Browse** to select the desired drive and directory. **Tips** • When you select an extended storage location and save your changes, Security Manager checks that it can be accessed and that it has write permissions. The primary storage location is used as a reference, and any data that exists in the primary storage location that does not exist in the extended storage location is copied to the extended storage location. Any data that already exists in the extended storage location is not evaluated and is left untouched, although it can be deleted later to make room for new data. • If you change the extended data store location, you cannot query events that exist only in the previous extended data store location (that is, you cannot query events that have already be removed from the primary location). If you want to preserve these events, copy the data from the old location to the new location.
Extended Data Store Disk Size	The amount of space you want to allocate to the extended event storage location, in gigabytes (GB). Events are incrementally deleted (rotated out) from the extended store when it becomes 90% full. The size must be equal to or larger than the primary event data storage location. You can see a visual representation of the amount of space currently used for event data. Open the Event Viewer (**Launch > Event Viewer**), then from Event Viewer, select **Views > Show Event Store Disk Usage**.

Table 11-16 Event Management Page (continued)

Element	Description
Error Notification Email IDs	The email addresses that should receive notifications if problems arise with the use of the extended storage location. Separate multiple addresses with commas. For notifications to be sent successfully, you must also configure an SMTP server as described in Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27. The message indicates the problem, cause, and recommended action. For example, you get notifications if the extended storage is chronically unreachable, if data copy fails repeatedly, or if a partition was deleted from the primary storage area before it could be copied to the extended storage area (which might happen if the storage is chronically unreachable or if there are persistent copy problems).
Syslogs for Failover Devices	
Process Syslogs from Failover Standby Device	Enables or disables processing of syslog messages from the standby ASA. When enabled, syslog messages generated by the standby or failover ASA will be displayed in the Device Identifier column in the Event Monitoring window. **Note** By default, the processing of syslog messages from the standby ASA is disabled.
Syslog Relay Service **Note** Staring with version 4.13, Cisco Security Manager supports syslogs over IPv6 in Event Viewer but the Syslog Relay Service will not be supported for syslogs over IPv6.	
Enable Syslog Relay Service	Enables or disables the Syslog Relay Service. Select the Enable Syslog Relay Service check box to enable the fields required for configuring the Syslog Relay Service.

Table 11-16 Event Management Page (continued)

Element	Description
Syslog Relay Capture Port	Specifies the UDP port on which the Syslog Relay Service listens for syslogs. The default is 514. If the Syslog Relay Service is enabled, devices must send syslogs to the Syslog Relay Capture Port so that they can be forwarded to the local collector and remote collectors. If the Syslog Relay Service is turned off, devices should send syslogs to the Event Syslog Capture Port. **Note** The Syslog Relay Capture Port and the Event Syslog Capture Port cannot be the same. When enabling the Syslog Relay Service, if devices are currently configured to send syslogs to the Event Syslog Capture Port, you should instead use that port number for the Syslog Relay Capture Port and then change the Event Syslog Capture Port to something else. You must ensure that the Security Manager server, and intervening firewalls, allow incoming traffic on this port for Security Manager to collect the events. Managed devices must be configured to send syslog information to this port on the Security Manager server. **Tip** If you change this port, you must also change the Syslog Servers policies for all ASA and FWSM devices and security contexts that send events to Security Manager. For more information, see Syslog Servers Page, page 54-27.
Relay to Local Event Collector	Enables or disables syslog relay for the local event collector.
Relay to Remote Collector 1	Enables or disables syslog relay for Remote Collector 1.
Collector 1 IP address	Specifies the IP address to which syslogs should be sent for Remote Collector 1.
Collector 1 Syslog Capture Port	Specifies the UDP port on which Remote Collector 1 is listening for relayed syslogs.
Relay to Remote Collector 2	Enables or disables syslog relay for Remote Collector 2.
Collector 2 IP address	Specifies the IP address to which syslogs should be sent for Remote Collector 2.
Collector 2 Syslog Capture Port	Specifies the UDP port on which Remote Collector 2 is listening for relayed syslogs.

Table 11-16 Event Management Page (continued)

Element	Description
Device Filter	You can filter the devices for which syslogs should be relayed for a specific collector. Using this feature, you can configure syslogs for one set of devices to go to one collector and syslogs for a different set of devices to go to another collector: 1. Select the tab (Local Collector, Remote Collector 1, or Remote Collector 2) for which you want to filter devices. 2. To specify the devices for which you want relay syslogs for this collector, select **Permit Relay**. If instead you want to specify the devices for which you want to disable syslog relay for this collector, clear the **Permit Relay** check box. If the Permit Relay check box is **not** selected, then syslogs for the devices you add to the filter will not be relayed; however, syslogs for all other devices will be relayed. **Note** For each enabled collector, syslog relays from all devices are enabled by default. **Note** When adding a cluster to the filter list, the IP addresses for the cluster management pool will be included as part of filter configuration. 3. Select the devices or device groups from the Available Devices list that you want to add to the filter and click >> to move them to the Selected Devices list. For more information on selecting devices, see Using Selectors, page 1-45. 4. To add a device that is not managed in Security Manager, enter the IP address of the device in the Add Special Device field and then click the bottom >> to move the device to the Selected Devices list.
Restart	Restarts the Syslog Relay Service.
CPU Throttle Settings	Opens the CPU Throttling Policy dialog box in which you can control the CPU load used by the syslog relay service. For more information, see CPU Throttling Policy Dialog Box, page 11-33.
View Statistics	Opens the Syslog Relay Statistics dialog box in which you can see the average CPU and memory usage of the syslog relay service process as well as traffic rates for the different collectors. For more information, see Syslog Relay Statistics Dialog Box, page 11-35.
Save button	Saves and applies changes. Most changes related to the Event Viewer settings require that the Event Manager service briefly stop and then restart. If you change whether the service is enabled, it stops or starts, as appropriate. You are shown a progress indicator. Changes to Syslog Relay Service settings require that the Syslog Relay Service briefly stop and then restart. If you change whether the service is enabled, it stops or starts, as appropriate.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Troubleshooting Syslog Relay Servers

If the Syslog Relay Service is enabled, devices must send syslogs to the Syslog Relay Capture Port so that they can be forwarded to the local collector and remote collectors. If the Syslog Relay Service is turned off, devices should send syslogs to the Event Syslog Capture Port.

Syslog Relay Servers act as an intermediate connection between device events and Security Manager Event Manager application. It receives device event packets and forwards them to the Local Collectors and Remote Collectors.

Device Management via IP

To manage devices in Security Manager via IP (using IPv4 or IPv6), the Device Management interface must have the appropriate IP information.

For example, see the following sample configuration.

```
!
interface Management1/1
management-only
nameif management
security-level 100
ip address 10.197.87.95 255.255.255.0
ipv6 address 2016::b2aa:77ff:fe7c:a068/64
ipv6 enable
```

In this configuration, the Device Management IP address has both IPv4 and IPv6 management addresses. So you can manage a device via IPv4 or IPv6.

Problem:

If a device is managed via IPv6 Management Address in Security Manager, the communication between Security Manager and the device would occur only via IPv6 address and not IPv4 address.

However, Event Syslog server still sends the Event Syslog packets only to IPv4 address, therefore in this scenario Security Manager cannot map the equivalent device for the received IPv4 Event Syslog packets.

When you add a filter device in Local Collector or Remote Collector for Syslog Relay Services in **Tools > Security Manager Administration > Event Management - Syslog Relay Service**, Security Manager tries to extract the device Management IPv4 address instead of the IPv6 management address.

However, there is no IPv4 Management Interface configured in the device. Therefore, Security Manager displays the following error:

```
Device selection – Ipv4 address not found for device(s)
```

Solution:

Go to **Device View > Policies > Interfaces**, to configure the device Management Interface with IPv4 address.

CPU Throttling Policy Dialog Box

Use the CPU Throttling Policy dialog box to specify settings for controlling the CPU load used by the syslog relay service.

After CPU throttling is enabled, if the syslog relay service's average CPU usage over the time period selected in the Average Max CPU Usage Time field is greater than the Maximum CPU Usage threshold, then CPU throttling will take place for the collectors specified in the Stop Forwarding To field for the time specified in the Stop Forwarding For field.

Note You can use the Syslog Relay Statistics dialog box to see the number of syslog packets dropped per collector due to the throttle policy (see Syslog Relay Statistics Dialog Box, page 11-35).

Navigation Path

Click **Tools > Security Manager Administration**, select **Event Management** from the table of contents, and then click **CPU Throttle Settings**.

Field Reference

Table 11-17 CPU Throttling Policy Dialog Box

Element	Description
Enable CPU Throttling	Whether to enable throttling for the syslog relay service. CPU throttling for the syslog relay service is disabled by default.
Maximum CPU Usage	Specify the maximum CPU usage for the syslog relay service as a percentage of total CPU capacity. This is threshold at which CPU throttling will be initiated.
Average Max CPU Usage Time (Minutes)	Specifies the time in minutes for which CPU usage by the syslog relay service is calculated. Options are 1 minute, 5 minutes, and 15 minutes. This average is compared to the Maximum CPU Usage value to determine whether throttling should take place.
Stop Forwarding To	Specify the collectors for which you want to stop forwarding syslogs when throttling is engaged.
Stop Forwarding For	Specify how long, in minutes, throttling should be enabled when the threshold is hit. After the specified interval of time has elapsed, if the CPU usage is still above the Maximum CPU Usage threshold, throttling will remain in effect.
Enable Email Notifications	Whether to send email notifications when the syslog relay service enters or exits throttle mode. Email notifications are disabled by default. For the e-mails to be sent, you must configure an SMTP server as described in Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.
Notification Email IDs	Enter one or more valid addresses in the Notification Email IDs field; separate multiple addresses with commas.

Table 11-17 CPU Throttling Policy Dialog Box (continued)

Element	Description
Send Email	Specify how often to send notification emails: • **Every time**—Select this option to have a notification sent every time the syslog relay service enters or exits throttle mode. If the CPU usage is still above the Maximum CPU Usage threshold after the Stop Forwarding For timer has elapsed, throttling will remain in effect and an additional notification will be sent. • **Every**—Select this option to have at most one notification sent when the syslog relay service enters or exits throttle mode during a specific period of time. If you select this option, specify the time period by entering the number of minutes or hours and then selecting the corresponding option (Min/Hour) from the drop-down list.

Syslog Relay Statistics Dialog Box

Use the Syslog Relay Statistics dialog box to view the average CPU and memory usage of the syslog relay service process as well as traffic rates for the different collectors.

Navigation Path

Click **Tools > Security Manager Administration**, select **Event Management** from the table of contents, and then click **View Statistics**.

Field Reference

Table 11-18 Syslog Relay Statistics Dialog Box

Element	Description
Log Relay Service	
Memory Usage Average (last 1 min.)	Shows the amount of memory used by the syslog relay service on average over the last minute.
CPU Usage Average (last 1 min.)	Shows the percentage of CPU capacity used by the syslog relay service on average over the last minute. **Tip** If the average CPU usage is too high, you might consider enabling CPU throttling for the syslog relay service (see CPU Throttling Policy Dialog Box, page 11-33).
Total sylog packets received	Shows the total number of syslog packets that have been received by the syslog relay service since the service was started.
Average syslog received per second since start	Shows the average number of syslog packets that have been received by the syslog relay service per second since the service was started.
Average syslog received per second for last 1 minute	Shows the average number of syslog packets that have been received by the syslog relay service per second over the last minute.
Average syslog received per second for last 5 minute	Shows the average number of syslog packets that have been received by the syslog relay service per second over the last five minutes.

Table 11-18 Syslog Relay Statistics Dialog Box (continued)

Element	Description
Average syslog received per second for last 15 minute	Shows the average number of syslog packets that have been received by the syslog relay service per second over the last fifteen minutes.
Period (in mins.) for which throttle policy is active	Shows how long, in minutes, the CPU throttle policy for the syslog relay service has been active. For more information, see CPU Throttling Policy Dialog Box, page 11-33.
Local Collector/Remote Collector 1/Remote Collector 2	
Total syslog packets sent successfully	Shows the total number of syslog packets that have been sent by the syslog relay service since the service was started.
Total syslog packets dropped (filter policy)	Shows the total number of syslog packets that have been dropped by the syslog relay service in accordance with the defined filter policy since the service was started.
Total syslog packets dropped (throttle policy)	Shows the total number of syslog packets that have been dropped by the syslog relay service in accordance with the throttle policy since the service was started.
Total syslog packets failed during transmit	Shows the total number of syslog packets that were not able to be forwarded by the syslog relay service since the service was started.
Average syslog sent per second since start	Shows the average number of syslog packets that have been sent by the syslog relay service per second since the service was started.
Average syslog sent per second for last 1 minute	Shows the average number of syslog packets that have been sent by the syslog relay service per second over the last minute.
Average syslog sent per second for last 5 minute	Shows the average number of syslog packets that have been sent by the syslog relay service per second over the last five minutes.
Average syslog sent per second for last 15 minute	Shows the average number of syslog packets that have been sent by the syslog relay service per second over the last fifteen minutes.
Refresh	Refreshes the statistics displayed on the Syslog Relay Statistics dialog box.

Health and Performance Monitor Page

Use the Health and Performance Monitor page of the Security Manager Administration window to enable network-wide health and performance monitoring. The Health and Performance Monitor (HPM) is a stand-alone application that lets you monitor key health and performance data for ASA devices, IPS devices, and VPN services by providing network-level visibility into device status and traffic information.

Tip If you get a message that the application is unavailable when you attempt to launch the Health and Performance Monitor, but the **Enable Health and Performance Monitor** option is selected on this page, try restarting Health and Performance Monitoring. First, deselect the Enable option and click Save. Wait for the service to stop. Then, select the Enable option, click Save, and wait for the service to finish restarting. You can then try opening the HPM application again.

Navigation Path

Click **Tools > Security Manager Administration** and select **Health and Performance Monitor** from the table of contents.

Field Reference

Table 11-19 Health and Performance Monitor Page

Element	Description
Enable Health and Performance Monitor	Lets you enable or disable the Health and Performance Monitoring service, which allows Security Manager to collect event information. If you disable this feature, you cannot use the HPM application. **Tip** If you change this setting and click **Save**, you are prompted to confirm that you want to start or stop the Health and Performance Monitoring service. If you click Yes, the service is started or stopped immediately, and you are shown a progress indicator and told when the change is completed. Wait until the status change is completed before continuing.
OOB Notification Settings **Note** To receive email notifications make sure that an SMTP server has been configured on the Security Manager Server. For more information, see, Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27 Cisco Security Manager considers an out-of-band (OOB) change to be any change made to a device manually or outside of Security Manager control, for example, by logging into the (monitored) device directly and entering configuration commands through the CLI. For devices monitored by the HPM application, Cisco Security Manager monitors the OOB changes, detected by the HPM periodically. If any out-of-band changes are detected, HPM generates an alert displayed on the **Device Status View** page and sends an email, to the configured recipients. **Note** If a Cisco Security Manager restart occurs during the update time, after an OOB change is detected and an email notification has already been sent, the same email maybe sent again after Cisco Security Manager starts up.	
Enable OOB Email Notification	Lets you enable or disable email notifications for Out of Band changes. **Note** When the email notification is disabled, only an alert is displayed on the **Device Status View** page. **Note** When HPM detects the OOB change and syncs with the Configuration Manager, a separate email alert notification is sent for each device being monitored. To prevent duplication, emails sent for each OOB change are tracked and stored in a file, once in 5 minutes. **Tip** The default tracking time is set as 5 minutes in the Cisco Security Manager properties file. You can update this as needed.
Recipient E-mail(s)	Specify the recipients who must be notified of the OOB change.

Table 11-19 Health and Performance Monitor Page (continued)

Element	Description
Save button	Saves and applies changes. Most changes require that the Health and Performance Monitoring service briefly stop and then restart. If you change whether the service is enabled, it stops or starts, as appropriate. You are shown a progress indicator.
Reset button	Resets changes to the last saved values.

Report Manager Page

Use the Report Manager page of the Security Manager Administration window to enable or disable the Report Manager feature in Cisco Security Manager. Report Manager is a stand-alone application that lets you view security and usage reports for devices and remote access IPsec and SSL VPNs.

Navigation Path

Click **Tools > Security Manager Administration** and select **Report Manager** from the table of contents.

Field Reference

Table 11-20 Health and Performance Monitoring Page

Element	Description
Enable Report Manager	Lets you enable or disable the Report Manager service. If you disable this feature, you cannot use the Report Manager application. **Tip** If you change this setting and click **Save**, you are prompted to confirm that you want to start or stop the Report Manager service. If you click Yes, the service is started or stopped immediately, and you are shown a progress indicator and told when the change is completed. Wait until the status change is completed before continuing.
Save button	Saves and applies changes. If you change whether the service is enabled, it stops or starts, as appropriate. You are shown a progress indicator.
Reset button	Resets changes to the last saved values.

Identity Settings Page

Use the Identity Settings page to configure the Active Directory (AD) server group to use for a NetBIOS domain for use with identity-aware firewall policies on ASA devices. These settings enable you to use the Find feature when selecting users or user groups for identity-aware policies or identity user group policy objects.

Tip You can also add entries by configuring the Identity Options policy on an ASA. When you save the policy, you are asked if you want to update the identity settings administrative page. Keep in mind that you can have a single domain-to-AD server match on the settings page, whereas you can configure different ASAs to use different server groups for a domain. Username lookup always selects the AD servers defined in the identity settings administrative page, regardless of what server group is configured for the individual ASA that you are configuring.

Navigation Path

Select **Tools > Security Manager Administration** and select **Identity Settings** from the table of contents.

Related Topics

- Creating Identity User Group Objects, page 13-19
- Selecting Identity Users in Policies, page 13-21

Field Reference

Table 11-21 Identity Settings Page

Element	Description
Domain-AD Server Group Mapping table.	Each row in the table defines the Active Directory (AD) server group to use for a NetBIOS domain for use with identity-aware firewall policies on ASA devices. • To add an entry, click the **Add Row (+)** button and fill in the Add AD Domain Server dialog box. See Domain AD Server Dialog Box, page 13-10. You need to enter the domain name and select the AAA server group object that specifies the LDAP AD servers. • To edit an entry, select it and click the **Edit Row (pencil)** button. • To delete an entry, select it and click the **Delete Row (trash can)** button. • To test whether Security Manager can successfully contact the servers defined in a server group, select the row and click **Test.**

Table 11-21 Identity Settings Page (continued)

Element	Description
Default Domain	The NetBIOS domain to use when you do not type in a domain when specifying a user or group name in a firewall policy or an identity user group policy object. The default is LOCAL, which means the name is defined on the ASA itself, either as a local user or as a VPN user who was authenticated by a means other than an LDAP server group associated with a domain name. Other than LOCAL, only domains configured in the Domain-AD Server Group Mapping table appear in this list. **Tip** This setting is not related to the default domain configured on the ASA using the user-identity default-domain command. This setting is a convenience setting to allow you to type in usernames without always having to include the domain name. Select the domain for which you will most often type user names.
Route query via	When you use the Find feature while selecting users or user groups, Security Manager must query the AD server. Select whether the query comes from the Security Manager client (the workstation on which you are running the client) or the server. By default, LDAP queries come from the client.
For user strings without domain	If you select something other than LOCAL for the default domain, how to handle username or user group names that you type in without a domain name: • **Auto determine user/user-group from AD**—Check the AD server associated with the default domain to determine whether the name is for a user or user group, and add the appropriate string: Default-Domain\user or Default-Domain\\user-group. If the name cannot be found, you must manually type in the domain name and the one or two \ characters to indicate whether the name is for a user or a group. • **Change it to Default-Domain/user**—Assume that typed in names are user names, not user group names, and add the default domain: Default-Domain\user. **Tip** When typing, if you precede the name with \ or \\, the default domain is automatically added. Thus, if you select the Change it to Default-Domain/user option, you can still enter group names without typing the domain by first entering \\.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Image Manager Page

Use the Image Manager page to control the administrative settings for Image Manager within Security Manager.

Navigation Path

Select **Tools > Security Manager Administration** and select **Image Manager** from the table of contents.

Field Reference

Table 11-22 Image Manager Page

Element	Description
Edit CCO Settings	Use the Edit CCO Settings link to quickly navigate to the CCO Settings page. For information on the CCO Settings page, see CCO Settings Page, page 11-4.
Purge Jobs Older Than	Enter the length of time in days to hold Image Manager jobs before purging them. The default is 365 days. Select **Purge Now** to immediately clear previous Image Manager job specifications.
Include Repository	If checked, the image repository is part of Security Manager backup. The default is to exclude images.
Save button	Saves your changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

IP Intelligence Settings Page

Use the IP Intelligence Settings page to control the administrative settings for the IP Intelligence features within Security Manager.

Navigation Path

Select **Tools > Security Manager Administration** and select **IP Intelligence Settings** from the table of contents.

Field Reference

Table 11-23 IP Intelligence Settings Page

Element	Description
Edit CCO Settings	Credentials for connecting to Cisco.com are required for automatic updates of the GeoIP database. You can use the **Edit CCO Settings** link to quickly navigate to the CCO Settings page where these credentials are configured. You can also configure settings for a proxy server on the CCO Settings page. For information on the CCO Settings page, see CCO Settings Page, page 11-4.

Table 11-23 IP Intelligence Settings Page (continued)

Element	Description
Reverse DNS (FQDN)	
Enable Reverse DNS (FQDN) Lookup Service	Whether to enable or disable the Reverse DNS (FQDN) lookup service. Enable this service if you want to be able to determine the fully qualified domain name (FQDN) for an IPv4 address using the IP Intelligence tool.
Use CSM Server's DNS Server	Select this option to use the DNS server defined on the Cisco Security Manager server for reverse DNS lookup requests.
Use custom DNS servers	Select this option to manually specify the DNS servers to use for reverse DNS lookup requests. You can enter up to three DNS server addresses in the fields provided. **Note** Security Manager does not support the use of external DNS servers configured inside of a virtual machine.
Enable Load Balancing	Whether to distribute reverse DNS lookup requests amongst the DNS servers when multiple DNS servers are available.
Default Blocking Ranges	Lists the IP address ranges that are excluded from Reverse DNS lookup by default: 0.0.0.0, 255.255.255.255, 127.0.0.1, 169.254.0.0-169.254.255.255, 224.0.0.0-239.255.255.255
User-defined Blocking Ranges	Specifies additional IP addresses or address ranges that should be excluded from reverse DNS lookup requests. Click the Edit (pencil) button to open the Edit IPv4 Blocking Range Addresses dialog box in which you can specify the IPv4 addresses or address ranges to be excluded. Separate multiple entries using a comma ",".
GeoIP	
Enable GeoIP Lookup Service	Whether to enable or disable the GeoIP lookup service. Enable this service if you want to be able to retrieve geographic location information for an IPv4 address using the IP Intelligence tool. **Note** You will need to download the geographic location database from Cisco.com before GeoIP information will be included in the IP intelligence data. You will also need to download the geographic location database from Cisco.com after restoring the Security Manager database from a backup. Beginning with version 4.9, Security Manager mandates you to read and accept the End User License Agreement (EULA) before you can proceed to downloading updates from cisco.com.

GeoIP Manual Upload

Use the GeoIP Manual Upload fields to update the geographic location database in Security Manager using a MaxMind GeoLite City update package downloaded from Cisco.com.

Note New update packages are made available on Cisco.com on a monthly basis.

Table 11-23 IP Intelligence Settings Page (continued)

Element	Description
GeoIP Database Artifact Location	Click **Browse** and then navigate to and select the MaxMind GeoLite City update package that you downloaded from Cisco.com. Then, click **Upload** to upload the selected database to Cisco Security Manager. **Note** Geolocation updates obtained directly from MaxMind or any other source are not supported in Cisco Security Manager.
GeoIP Maxmind Database Update Settings MaxMind GeoLite City update packages are updated monthly on Cisco.com. Use the GeoIP Maxmind Database Update Settings to download an update package automatically from Cisco.com and to configure scheduled updates. **Note** Credentials for connecting to Cisco.com are required for automatic updates of the geographic location database. You can use the **Edit CCO Settings** link to quickly navigate to the CCO Settings page where these credentials are configured. For information on the CCO Settings page, see CCO Settings Page, page 11-4.	
Run immediate database update	Click **Update Now** to update the geographic location database in Security Manager using the latest update package from Cisco.com.
Enable scheduled update	Whether to enable or disable automatic updates of the geographic location database on a regular schedule. After enabling scheduled updates, click **Edit Settings** to specify the schedule for when the update should take place. Using the Weekly option, you can specify the days of the week on which the automatic update should take place. Using the Monthly option, you can specify the day of the month on which the automatic update should take place. For either option, you can specify the time of day that the update should take place. **Tip** The geographic location database is updated by MaxMind on the first Tuesday of every month. New update packages are typically available on Cisco.com approximately one week after they are issued by MaxMind. We recommend configuring your update schedule to occur monthly on day 15 or later. However, you can schedule the updates to take place more frequently if you want to ensure that the updated database is made available in Security Manager as close to the time it is available on Cisco.com as possible. **Note** Beginning with version 4.9, Security Manager mandates you to read and accept the End User License Agreement (EULA) before you can proceed to downloading updates from cisco.com.
Default Blocking Ranges	Lists the IP address ranges that are excluded from GeoIP lookup by default: 0.0.0.0, 255.255.255.255, 127.0.0.1, 10.0.0.0-10.255.255.255, 169.254.0.0-169.254.255.255, 172.16.0.0-172.31.255.255, 192.168.0.0-192.168.255.255, 224.0.0.0-239.255.255.255

Table 11-23 IP Intelligence Settings Page (continued)

Element	Description
User-defined Blocking Ranges	Specifies additional IP addresses or address ranges that should be excluded from GeoIP lookup requests. Click the Edit (pencil) button to open the Edit IPv4 Blocking Range Addresses dialog box in which you can specify the IPv4 addresses or address ranges to be excluded. Separate multiple entries using a comma ",".
Whois	
Enable Whois Lookup Service	Whether to enable or disable the Whois lookup service. Enable this service if you want to be able to retrieve WHOIS information for an IPv4 address using the IP Intelligence tool.
Enable External Proxy	Whether to enable or disable use of an external proxy for Whois requests. Proxy server configuration is specified on the CCO Settings page. **Tip** You can use the **Edit CCO Settings** link to quickly navigate to the CCO Settings page where the proxy server settings are configured. For information on the CCO Settings page, see CCO Settings Page, page 11-4.
Default Blocking Ranges	Lists the IP address ranges that are excluded from Whois lookup by default: 0.0.0.0, 255.255.255.255, 127.0.0.1, 10.0.0.0-10.255.255.255, 169.254.0.0-169.254.255.255, 172.16.0.0-172.31.255.255, 192.168.0.0-192.168.255.255, 224.0.0.0-239.255.255.255
User-defined Blocking Ranges	Specifies additional IP addresses or address ranges that should be excluded from Whois lookup requests. Click the Edit (pencil) button to open the Edit IPv4 Blocking Range Addresses dialog box in which you can specify the IPv4 addresses or address ranges to be excluded. Separate multiple entries using a comma ",".
View Statistics	Opens the IP Intelligence Statistics dialog box which shows statistics for the IP Intelligence feature. Information provided in the IP Intelligence Statistics dialog box includes: • Average number of IP Intelligence lookup requests during the last 5 minutes and last 15 minutes • Average lookup time for all IP Intelligence service requests • Average lookup times for each individual service that is currently enabled • Counts of the number of lookups, both successful and failed, for each individual service that is currently enabled • Cache hit ratios for each individual service that is currently enabled • Upload information for GeoIP updates: last update time, status, and version information for the update Click **Refresh** to update the data in the IP Intelligence Statistics dialog box.
Save button	Saves your changes.

Table 11-23 IP Intelligence Settings Page (continued)

Element	Description
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Eventing Notification Settings Page

Use the Eventing Notification Settings page to receive email notifications for IPS events and critical ASA events. You can configure the time interval at which you want to receive the email notifications.

The events are sent in the form of .CSV files in a .zip file format. By default, email notification is disabled. When you enable email notification, only the notification for IPS events is enabled. To receive email notification for critical events you must enable the additional settings for critical events.

Note For notifications to be sent successfully by Security Manager, you must configure an SMTP server, as described in Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.

Tip You can also use the Security Manager Event Viewer application or the Dashboard to view and monitor all events.

Navigation Path

Select **Tools > Security Manager Administration** and select **Eventing Notification Settings** from the table of contents.

Field Reference

Table 11-24 Eventing Notification Settings Page

Element	Description
Enable Eventing Email Notification	Select to enable notification of IPS events through email.
Notification Interval (15 - 60 minutes)	Enter the interval at which you want Security Manager to send email notifications for IPS events or critical events. **Note** If Security Manager receives more than 50000 events during the configured time interval, only the first 50000 events are selected and sent through email.
Notification Settings (IPS)	
Email IDs (for IPS Events)	Enter one or more email addresses (comma separated).
Select Severity of Events	The severity level that the IPS signature reports: High, Medium, Low, or Informational. By default, High and Medium severity are selected.
Notification Content	Whether to send summarized or detailed notifications through email. If you select Detailed Notifications, select the fields that you want information for, in the email notification. Some fields are selected by default.

Table 11-24 Eventing Notification Settings Page (continued)

Element	Description
Fields	
Event ID	A unique sequential number for each event, assigned internally.
Severity	The Firewall or IPS severity values.
Device	The source of the event; usually the device ID. A device identified as Not Available has been deleted from the Security Manager inventory.
Application Name	The name of the application originating the event.
Receive Time	The time the event was received by Security Manager.
Event Time	The time the event was generated by the device.
Sensor Local Time	The local time of the sensor where the event originated.
Sig ID	The Sig ID value is used by the alert originator to identify the activity. It identifies the pre-defined signature defined for this activity.
Sub Sig. ID	Identifies the unique numerical value assigned to the sub signature. The Sub Sig. Id identifies a more granular version of a broad signature.
Sig. Name	The name assigned to the signature.
Sig. Details	The details of the reported signature that was triggered and resulted in the generation of the alert.
Sig. Version	The version of the signature definition used to generate an alert.
Attacker IP	The IP address of the host that sent the offending packet.
Attacker Port	The port used by the attacker host. This is the port from which the offending packet originated.
Attacker Locality	Identifies whether the attacker address is located inside or outside of a given network, as specified by the intrusion detection device's configuration.
Victim IP	The IP address of the host being attacked.
Victim Port	The port of the host being attacked (the recipient of the offending packet). This is the port to which the offending packet was sent.
Victim OS	The OS of the host being attacked.
Victim Locality	Identifies whether the target address is located inside or outside of a given network, as specified by the intrusion detection device's configuration.
Summary Count	Specifies that this is a summary alert, representing one or more alerts with common characteristics. The numeric value indicates the number of times the signature fired since the last summary alert with a matching initialAlert attribute value.
Initial Alert	This field applies to a summary alert, representing one or more alerts with common characteristics. The value of InitialAlert provides the event ID of the last nonsummary evIdsAlert with the same characteristic (sigid/ subsigid).
Summary Type	Defines the common characteristics of all alerts in a summary alert.

Table 11-24 Eventing Notification Settings Page (continued)

Element	Description
Is Final	Applies to a summary alert, representing one or more alerts with common characteristics. It indicates whether this is the last event alert containing the same value in the initialAlert attribute.
Interface	Name of the IPS interface.
VLAN	The VLAN number associated with packets involved in the activity that triggered the alert.
Virtual Sensor	The name of the virtual sensor associated with the event.
Action Taken	The action performed on the flow. For example: Terminated or denied.
Alert Details	The details regarding the alerts.
Risk Rating	A value that represents the calculated risk associated with the event.
Threat Rating	The threat rating of the event, if any.
Reputation	The attacker's reputation score in the range -10.0 to +10.0. A lower (more negative) score indicates a greater likelihood that the host is malicious.
Reputation Details	Deny Attacker: Whether a deny-attacker action occurred (or would have occurred) because an internal override was exceeded due to the calculated risk rating: true or false.
Protocol	The Level-3 or Level-4 protocol.
Notification Settings (Critical Events only)	
Enable Email IDs	Whether to send email notification for critical events. If you select this option, also enter one or more email addresses (comma separated).
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

IPS Updates Page

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

Use the IPS Updates page to perform administrative tasks associated with keeping your sensors up to date with regard to signatures, minor version updates, and service packs. You can use the IPS Updates page to:

- Monitor update status.
- Check the availability of updates and download them.
- Configure an IPS update server.
- Configure automatic update settings.

Note Beginning with Security Manager version 4.9, only the latest sensor and signature packages for IPS will be available for download from CCO. The older packages will not be available for download from CCO.

Tips

- To apply IPS updates manually, select **Tools > Apply IPS Update**. For more information, see Manually Applying IPS Updates, page 44-7.
- If you later decide that you did not want to apply a signature update, you can revert to the previous update level by selecting the Signatures policy on the device, clicking the **View Update Level** button, and clicking **Revert**.

Beginning with version 4.4, Security Manager has a certificate trust management feature. This feature helps you with improved handling of Cisco.com certificates. For detailed documentation of this feature, refer to Certificate Trust Management, page 10-18.

Navigation Path

Select **Tools > Security Manager Administration** and select **IPS Updates** from the table of contents.

Related Topics

- Configuring the IPS Update Server, page 44-4
- Checking for IPS Updates and Downloading Them, page 44-5
- Automating IPS Updates, page 44-6
- Selecting a Signature Category for Cisco IOS IPS, page 45-6

Field Reference

Table 11-25 IPS Updates Page

Element	Description
Update Status group Refresh button	Displays the following items. Click **Refresh** to update the information. • Latest Available—The most recent signature and sensor update available on Cisco.com or the local HTTP server when you last checked for updates. • Latest Downloaded—The most recent signature and sensor update downloaded to Security Manager. • Latest Applied—The most recent signature and sensor update applied to any device in Security Manager. • Latest Deployed—The most recent signature and sensor update deployed to any device in Security Manager. • Last Check On—The time that the last check of Cisco.com was performed. • Last Download On—The time that the last update was downloaded to Security Manager. • Last Deployed On—The time that the last update was deployed to any of the devices.

Table 11-25 IPS Updates Page (continued)

Element	Description
Check for Updates button Download Latest Updates button	These buttons check for updates, or download signature and sensor updates that have not already been downloaded to the Security Manager server, from the IPS Update server. You must configure an IPS Update server before checking for updates or downloading them (click **Edit Settings** in the Update Server group). When you click one of these buttons, a dialog box opens to display the results of the operation. Security Manager logs into the IPS Update server, checks for updates, and downloads them if you clicked the Download button. If a Cisco.com download fails, ensure that the account you are using has applied for eligibility to download strong encryption software. For details, see the description of User Name in Edit Update Server Settings Dialog Box, page 11-52. **Tip** If you configure a server, and then try to check for updates, and you are told you did not configure a server, click **Save** at the bottom of the page and try again. **Note** Beginning with version 4.9, Security Manager mandates you to read and accept the End User License Agreement (EULA) before you can proceed to downloading updates from cisco.com.
Update Server group	Displays the settings used to access Cisco.com or the local server that contains the IPS update packages. The fields indicate whether the update server is Cisco.com or a locally-configured HTTP server, the name of the local server if you are using one, the user account for logging into the server, and the name of the proxy server, if any. To configure or change the IPS Update server, click **Edit Settings** to open the Edit Update Server Settings dialog box (see Edit Update Server Settings Dialog Box, page 11-52). For more information, see Configuring the IPS Update Server, page 44-4 Beginning with version 4.4, Security Manager has a certificate trust management feature. This feature helps you with improved handling of Cisco.com certificates. For detailed documentation of this feature, refer to Certificate Trust Management, page 10-18.
Signature Filter Settings group	Enables you to download IPS signature updates selectively. Click **Edit Settings** to open the Edit Signature Download Filter Settings dialog box (see Edit Signature Download Filter Settings Dialog Box, page 11-55).
Auto Update Settings group	Contains the settings specific to automatic updates. For more information, see Automating IPS Updates, page 44-6.

Table 11-25 IPS Updates Page (continued)

Element	Description
Auto Update Mode	Establishes whether, and to what extent, automatic updates are performed. Contains the following options: • Download, Apply, and Deploy Updates • Disable Auto Update • Check for Updates • Download Updates • Download and Apply Updates By default, auto update is disabled. The other options are a combination of one or more of the following options: • Check for Updates—Security Manager contacts the IPS Update server to check if an update is available and sends e-mail if e-mail notification is configured. No files are downloaded. • Download Updates—Security Manager downloads the latest updates from the IPS Update server, and sends e-mail notification if e-mail notification is configured. • Apply Updates—Security Manager modifies the configuration of the devices selected in the Apply Update To list based on the downloaded update packages. You have to separately deploy these updates unless you also select Deploy Updates. • Deploy Updates—Security Manager starts a deployment job to send the applicable update packages to the devices selected in the Apply Update To list. The device must have the required license for a signature update to be successful.
Update Schedule Edit Update Schedule button	The schedule for the actions selected in the Auto Update Mode field. To change the schedule, click **Edit Update Schedule** and define the schedule in the Edit IPS Updates Schedule dialog box. You can specify that Security Manager perform the updates based on hourly, daily, weekly, or monthly schedules, or specify a one-time event. When entering the start time, use the 24-hour clock and the *hh:mm* format. **Note** If you schedule an update to occur in less than 10 minutes from your Security Manager server time, the "Next Update" field will show tomorrow's date and the job will run accordingly. This is a safety feature designed to guarantee the first occurrence to run. **Tip** Cisco recommends scheduling automatic downloads during off hours so that they do not conflict with other user operations, such as device discovery. **Tip** Cisco recommends using an account other than the admin account for routine user operations.

Table 11-25 IPS Updates Page (continued)

Element	Description
Notify Email	The e-mail address to which notifications of automatic updates are sent. If you enter more than one address, separate the addresses with commas. A notification is sent when an update: • Is available for download. • Has been downloaded. • Has been downloaded and applied. • Has been downloaded, applied, and deployed.
Apply Update To Type Edit Row button Devices to be Auto Updated	The selector includes the IPS devices that have local signature policies and the shared signature policies that are defined in Security Manager. The columns in the selector indicate whether a local device policy or a shared policy is selected for these types of updates: • Signature—For auto updating the signature update level. • Minor—For minor updates and service packs. • S.P.—For service pack updates. For shared policies, a partial grey checked box indicates that some, but not all, of the devices that use the policy are selected. If you change the devices assigned to the shared policy between automatic update events, the shared policy is grayed out, and only the old assignments are shown on this page. After the update runs, the assignment list will be synchronized with the shared policy device assignments. To update the device list proactively prior to the next auto update run, select the policy and edit it (to select auto update settings), and the device assignment list will be corrected. **Note** Also for shared policies: You can select only the shared policy assigned to the default virtual sensor (vs0). If you attempt to select the shared policy for a different virtual sensor, your changes will not be applied, and you will not receive an error message. Use the Type field to toggle between viewing local and shared policies. Changing the view does not change your auto update selections. To select a local or shared policy for auto update, select it in the selector and click the **Edit Row** button below the selector. This opens the Edit Auto Update Settings dialog box, where you can select the types of updates for the policy. When you select any type of auto update for a policy, the affected devices are listed in the **Devices to be Auto Updated** list to the right of the selector.
Save button	Saves your changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Edit Update Server Settings Dialog Box

Use the Edit Update Server Settings dialog box to configure the server to use for obtaining IPS updates. If necessary, you can configure a proxy server for communicating with the update server.

Also, use the Edit Update Server Settings dialog box for certificate trust management. (Security Manager downloads IPS packages from Cisco.com over HTTPS, which uses certificates for establishing trust.) The certificate trust management feature on the Image Manager page is new in Security Manager 4.4. It will help you with improved handling of Cisco.com certificates for IPS package downloads:

- You can use it to view a certificate and use discretion in accepting it.
- After you accept a certificate, it is stored on your Security Manager server.
- You can see all your certificates in a summary table on the Image Manager page, and you can use that table to view or remove certificates.

Tip Please be sure to refer to "Retrieve Certificate" in the table below.

Navigation Path

Select **Tools > Security Manager Administration > IPS Updates** and click **Edit Settings** in the Update Server group.

Field Reference

Table 11-26 Edit Update Server Settings Dialog Box

Element	Description
Update From	Whether to get IPS updates from Cisco.com or from a local HTTP/HTTPS server. Your selection changes the fields on the dialog box. If you select local, you must configure an HTTP or HTTPS server to use as the IPS update server. **Caution** The default value for "Update From:" is "Local Server." You must choose "Cisco.com" to see certificate settings. Improper or incomplete certificate setting will prevent connectivity to Cisco.com, and all Cisco.com-related operations in this area will fail.
IP Address/ Host Name (Local server only.)	The hostname or IP address of the local IPS update web server.
Web Server Port (Local server only.)	The port number that your local server listens to for connection requests. The default is 80.

Table 11-26 Edit Update Server Settings Dialog Box (continued)

Element	Description
User Name	The username to log into the IPS update server. If you are configuring a local server that does not require a user login, leave this field blank. If you are specifying a Cisco.com username, the user account on Cisco.com must be eligible for downloading strong encryption software. If you are not certain that the account has the required permissions, use the account to log into Cisco.com and try to download an IPS update file (http://www.cisco.com/cgi-bin/tablebuild.pl/ips5-system). If the account does not have the appropriate permissions, you are prompted to read and accept the required conditions. If you meet the eligibility requirements, you can accept them. Otherwise, talk to your Cisco sales representative for help.
Password Confirm	The password for the specified username, entered in both fields. If you are configuring a local server that does not require a password, leave these fields blank.
Path to Update Files (Local server only.)	The path to the IPS update files location on your local server. For example, if update files can be accessed at http://local-server-ip:port/update_files_path/, then enter update_files_path in this field.
Connect Using HTTPS (Local server only.)	Whether to use SSL when connecting to the local IPS Update server.
Certificate Thumbprint	Displays the certificate thumbprint after it is calculated from the certificate on the local server.
Retrieve From Server	Used to connect to the local server specified in this dialog box, retrieve the certificate from the local server given, and calculate the certificate thumbprint, which is displayed in the Certificate Thumbprint field.
Contact URL	• When selected, "Image Meta-data Locator" is used. This is the URL on Cisco.com that is used to obtain meta-data information about images. Meta-data information consists of the images applicable to a particular product, name, size, checksum, and URL to download for each image. • When selected, "Other" is used. You can enter any valid HTTPS URL. This URL is intended primarily for the HTTPS URL to download the image as obtained from the meta-data information about the image. This URL may be different from the URL of the image meta-data locator described in the previous paragraph; the certificate may be different, as well. ⚠ **Caution** If you choose "Other," you need to explicitly add "https://dl.cisco.com" [without the quotation marks]: Enter it in the text field adjacent to the "Other" button. Failure to do this will prevent connectivity to Cisco.com, and all Cisco.com-related operations in this area will fail.

Table 11-26 Edit Update Server Settings Dialog Box (continued)

Element	Description
Retrieve Certificate	Used to connect to and retrieve the certificate from the selected '"Contact URL'." After retrieving the certificate it opens the Certificate Verification dialog, which along with a brief summary of the certificate, i.e., who the certificate is issued to, by whom, and the validity period of the certificate, gives you the following choices: • **View Certificate**—Opens the Certificate Viewer, where you can see all the details of the certificate: Certificate Authority, version, serial number, thumbprint, and other details. It shows the complete certificate chain information all the way up to the root issuing certificate Authority. • **Accept**—Accepts the certificate and adds it to the Cisco Security Manager. • **Reject**—Rejects the certificate and no action is taken. • **Cancel**—Closes the Certificate Verification dialog with no action taken.
Certificate	A table that displays, for each certificate in your Security Manager installation, Subject, Issued By, and Accepted By.
View	Opens the Certificate Viewer for a certificate selected in the Certificate table.
Remove	Removes a certificate selected in the Certificate table.
Proxy Server Group	
Enable Proxy Server	Whether a proxy server is needed to connect to Cisco.com or to your local server.
IP Address/ Host Name	The hostname or IP address of the proxy server. You can configure the proxy server to use basic, digest, NT LAN Manager (NTLM) V1, or NTLM V2 authentication. NTLM V2 is the most secure scheme.
Port	The port number that the proxy server listens to for connection requests. The default is 80.
User Name	The username to log into the proxy server. If the proxy server does not require a user login, leave this field blank.
Password Confirm	The password for the specified username, entered in both fields. If the proxy server does not require a password, leave these fields blank.

Edit Auto Update Settings Dialog Box

Use the Edit Auto Update Settings dialog box to configure the automatic update options for the device or policy selected in the Apply Update To table on the IPS Updates page. For information on configuring automatic updates, see Automating IPS Updates, page 44-6.

Navigation Path

Select a device or policy on in the Apply Update To table on the IPS Updates page (see IPS Updates Page, page 11-47) and click the Edit Row button.

Field Reference

Table 11-27 Edit Auto Update Settings Dialog Box

Element	Description
Auto Update (IPS sensors and shared policies only)	The type of sensor updates to apply to the selected devices or shared policies. You can apply both minor updates and service packs, service packs only, or select None to apply no sensor updates automatically.
Auto Update Signature Update Level	Whether to select the device or policy for automatic signature updates.

Edit Signature Download Filter Settings Dialog Box

The Edit Signature Download Filter Settings dialog box enables you to download IPS signature updates selectively. It applies both to the manual download and to the automated download.

Note Filtering does not apply to IPS sensor packages or to IPS engine packages; it applies to IPS signature packages only. All the available sensor packages on Cisco.com or on the local server will be downloaded as part of a signature download.

The benefits of selective download are reduced download time, reduced disk storage space, and faster troubleshooting because you can download only what you need.

There are four types of signature download available to you with the Edit Signature Download Filter Settings dialog box:

- No filter
- Download all signatures for engine versions starting with [choose E4, E3, E2, or E1]
- Download all signature versions starting with [enter a signature version such as 1000]
- Download a single signature version number [enter a signature number such as 1000]

The default signature configuration is to download all signatures for engine versions starting with E4.

Tip This default value is the same for a new installation of Security Manager 4.3 and for upgrades from previous versions.

Navigation Path

Select **Tools > Security Manager Administration** and then select **IPS Updates** from the table of contents; then click **Edit Settings** in the Signature Filter Settings group.

Related Topics

- Configuring the IPS Update Server, page 44-4
- Checking for IPS Updates and Downloading Them, page 44-5

- Automating IPS Updates, page 44-6

Field Reference

Table 11-28 Edit Signature Download Filter Settings Dialog Box

Element	Description
Filter Type: No filter	All available signatures for all available engines are downloaded.
Filter Type: Download all signatures for engine versions starting with	All available signatures for the engine that you select (E4, E3, E2, or E1) are downloaded.
Filter Type: Download all signature versions starting with	All available signatures starting with the ID that you enter are downloaded.
Filter Type: Download single signature version number	The single signature having the ID that you enter is downloaded.

ISE Settings Page

Use the ISE Settings page to configure communication between Cisco Security Manager and the Cisco Identity Services Engine (ISE) for use with TrustSec firewall policies.

Note Security Manager supports communications with only one ISE appliance/server for fetching and resolving security group names and tags.

Navigation Path

Select **Tools > Security Manager Administration** and select **ISE Settings** from the table of contents.

Related Topics

- Chapter 14, “Managing TrustSec Firewall Policies”
- Creating Security Group Objects, page 14-14
- Selecting Security Groups in Policies, page 14-16

Field Reference

Table 11-29 Identity Settings Page

Element	Description
Enable ISE feature	Whether to enable communication with the ISE.
Username	The username Security Manager should use to log on to the ISE.
Password	The password for the username.
ISE Server (IP Address/Hostname)	The DNS hostname or IP address of the ISE.
ISE Version	Beginning with version 4.18, Cisco Security Manager supports integration of only ISE version 2.3.

Table 11-29 Identity Settings Page (continued)

Element	Description
Test Connectivity	Click Test Connectivity to ensure that Security Manager can communicate with the ISE given the settings you have entered.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Licensing Page

Use the Licensing page to manage licenses for the Security Manager application and for IPS devices. For more information, see Managing IPS Licenses, page 44-1.

Navigation Path

Select **Tools > Security Manager Administration** and select **Licensing** from the table of contents.

Field Reference

Table 11-30 Licensing Page

Element	Description
CSM tab	The license settings for the Security Manager application. For a description of the fields on this tab, see CSM Tab, Licensing Page, page 11-57.
IPS tab	The license settings for IPS devices managed by Security Manager. For a description of the fields on this tab, see IPS Tab, Licensing Page, page 11-58.

CSM Tab, Licensing Page

Use the CSM tab on the Licensing page to view the list of installed Security Manager licenses and to install new licenses. For more information, see Installing Security Manager License Files, page 10-16.

Navigation Path

Select **Tools > Security Manager Administration**, select **Licensing** from the table of contents, and click **CSM**.

Field Reference

Table 11-31 CSM Tab, Licensing Page

Element	Description
License Information	Displays information about the license registered with the product: the edition, license type, expiration date, the number of licensed devices, the number of devices in use, and the percentage of the device count used.

Table 11-31 CSM Tab, Licensing Page (continued)

Element	Description
Install License	The list of installed licenses with their installation dates.
Install a License button	Click this button to install a license file. The dialog box that is opened includes links to Cisco.com, where you can obtain licenses if you have not already obtained them. You must copy license files to a local drive on the Security Manager server before you can install them.

IPS Tab, Licensing Page

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

Use the IPS tab on the Licensing page to view the list of installed IPS device licenses, to install new or updated licenses, or to redeploy licenses. The license list shows current licenses, unlicensed devices, devices with expired licenses, and devices with invalid licenses. You can also use the settings on this page to send a report of all those IPS devices whose license would expire within a specified number of days.

Navigation Path

Select **Tools > Security Manager Administration**, select **Licensing** from the table of contents, and click **IPS**.

Related Topics

- Updating IPS License Files, page 44-1
- Redeploying IPS License Files, page 44-2
- Automating IPS License File Updates, page 44-3
- License Update Status Details Dialog Box, page 11-62
- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49

Field Reference

Table 11-32 IPS Tab, Licensing Page

Element	Description
IPS License Table	Displays all the IPS devices in the device inventory and their license status as of the last time you refreshed the information. Click the **Refresh** button to obtain the latest information from the devices. Information includes the serial number for the device, which is used to register for licenses, the license status, and the expiration date of the license. The list shows not only current licenses, but also unlicensed devices, devices with expired licenses, and devices with invalid licenses. **Tip** The list does not include Cisco IOS IPS devices. You cannot use Security Manager to manage licenses for routers running IPS.
Update Selected via CCO button	Click this button to update the license file for the selected devices by connecting to Cisco.com and retrieving a new license. When you click this button, a dialog box opens listing devices that can be updated from Cisco.com, which might not be all the devices you selected. Click **OK** to perform the update. For successful updates, the updated file is automatically applied to the device. To successfully update the license using this method, you must have a Cisco.com support contract that includes the serial numbers of the selected devices. **Tip** The Cisco software license server (SWIFT) that contains the licenses might block requests from the same server for more than 9 licenses within a three minute period. Thus, you should select fewer than 9 devices at a time when performing manual license updates.
Redeploy Selected Licenses button	Click this button to redeploy licenses to the selected devices. Redeploying licenses might be necessary when you have obtained an updated license file and it was not applied to the device successfully during an automatic update. When you click this button, a dialog box opens listing devices whose licenses you are redeploying. Click **OK** to perform the update. For successful updates, the updated file is automatically applied to the device.
Update from License File button	Click this button to update licenses by selecting a license file from the Security Manager server. When you click this button, a dialog box opens where you can specify the license files. Click **Browse** to select the files, which must be on a local drive on the Security Manager server. When you click **OK**, the updated files are automatically applied to the devices.
Export As button	Select one or more IPS devices from the list and then click the **Export As** button to export their licenses to a Portable Document Format (PDF) or comma-separated values (CSV) file. You are prompted to select the folder on the Security Manager server and to specify a file name. If you do not select any device from the list, the licenses of all available devices are exported.

Table 11-32 IPS Tab, Licensing Page (continued)

Element	Description
Refresh License button	Click this button to refresh the data in the IPS license table for the selected devices. The updated information is retrieved from the device. If you do not select any devices, all devices are refreshed; this can take a long time depending on the number of devices listed.
Download and apply licenses Days before the expiration date.	Whether to automatically download IPS licenses from Cisco.com and apply them to the devices. To successfully configure automatic updates, you must have a Cisco.com support contract that includes the serial numbers of your IPS devices. If you select this option, also specify the number of days before the license expiration date for downloading and applying licenses. Security Manager evaluates only those devices that do not have licenses, have expired licenses, or have valid licenses within this number of days of expiration. Licenses are applied only if they are valid and either have an expiration date farther out than the current one, or that have different license information.
Discover devices daily at	If you select automatic license updates, the time of day when Security Manager should contact devices for current licenses status and evaluate whether there are devices that have licenses that will expire within the specified number of days. Cisco.com is contacted only if one or more device meets the expiration requirements.
Email License Update Results Email Notification	Whether to send email notifications of expiration alerts and license update job results. If you select this option, also enter one or more email addresses (comma separated).
Email License Expiration Status Email Notification	Whether to send a PDF report of those IPS devices whose license would expire within a specified number of days. If you select this option: • Enter the number of days (not more than 100) before the device license expiry date, by which you want Security Manager to send the PDF report. • Select the time and day for Security Manager to check the license expiry. • Enter one or more email addresses (comma separated) to which you want the License Expiration Status PDF report to be sent.
Save button	Saves your changes to the automatic license update and e-mail notification settings.

Verifying IPS Devices for License Update or Redeployment

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

When you select a device on the **Licensing > IPS** tab (see IPS Tab, Licensing Page, page 11-58) and try to update the license from Cisco.com (CCO) or redeploy the license, you are first shown a list of devices that will be updated. The name of the dialog box is based on the action you are taking:

- **Updating Licenses via CCO dialog box**—Review the IPS devices you selected to update from Cisco.com. The device list displays the IPS devices for which you can update the license from Cisco.com, which might not be all of the devices you selected.

  To successfully update the license using this method, you must have a Cisco.com support contract that includes the serial numbers of the selected devices.

Tip The Cisco software license server (SWIFT) that contains the licenses might block requests from the same server for more than 9 licenses within a three minute period. Thus, you should select fewer than 9 devices at a time when performing manual license updates.

- **Redeploying Licenses dialog box**—Review the IPS devices you selected for redeploying licenses. Before you can redeploy a license to a device, you must have already deployed the license. Security Manager uses the file already associated with the IPS device to redeploy the license.

When you click **OK**, the License Update Status Details dialog box opens so that you can view the status of the license redeployment task. See License Update Status Details Dialog Box, page 11-62.

Navigation Path

To open these dialog boxes, select one or more device on the **Tools > Security Manager Administration > Licensing > IPS** tab and click **Update Selected via CCO** or **Redeploy Selected Licenses**.

Selecting IPS License Files

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements.

If you select one or more devices on the **Tools > Security Manager Administration > Licensing > IPS** tab and click **Update from License File**, you are prompted to select the license file you want to use with the Updating Licenses from File dialog box.

You can store the license file on a local drive on the Security Manager server, and, beginning with Version 4.5 of Security Manager, you can store it on a local drive on a client.

Click **Browse** to select the license file. You can select multiple license files using Ctrl+click or a range of files using Shift+click.

Note If you installed the Security Manager client on a different machine than the one on which Security Manager server is installed, you can choose to select the license file from either the client machine or the server machine. If both the client and the server are installed on the same machine, Security Manager allows you to select the license file only from the server.

When you have selected the license files you want to use, click **OK** to apply them to the IPS devices.

Note If you want to store the license file on a client machine, you must select "Enable Client side file browser" on the Customize Desktop page at Tools > Security Manager Administration > Customize Desktop.

License Update Status Details Dialog Box

Use the License Update Status Details dialog box to view the status of an IPS license update task. This dialog box opens whenever you start an update task from the IPS tab of the Licensing page. For more information, see IPS Tab, Licensing Page, page 11-58.

Field Reference

Table 11-33 License Update Status Details Dialog Box

Element	Description
Progress bar	Indicates what percentage of the license update task on the current device has been completed.
Status	The current state of the update task.
Devices to be updated	The total number of devices being updated during this task.
Devices updated successfully	The number of devices updated without errors.
Devices updated with errors	The number of devices that generated errors during the update.
Device list	The devices that are being updated, including the device name, the status of the update, and summary information about the update. Select a device to see the messages generated during the update for that device in the message list below the summary list.
Messages list	The messages generated during the license update for the selected device. Select a message to see detailed information in the fields to the right of the list.
Description	Additional information about the message selected in the message list.
Action	The steps you should take to resolve the described problem.
Abort button	Aborts the license update task.

Logs Page

Use the Logs page to configure the default settings for the audit and operations logs. The audit log keeps a record of all state changes that occur in Security Manager.

Navigation Path

Select **Tools > Security Manager Administration** and select **Logs** from the table of contents.

Related Topics

- Using the Audit Report Window, page 10-21
- Understanding Audit Reports, page 10-19
- Generating the Audit Report, page 10-20
- Purging Audit Log Entries, page 10-23

Field Reference

Table 11-34 Logs Page

Element	Description
Manual Purging	
Keep Audit Log in DB For (days)	The maximum number of days of Audit Log entries that has to be stored in the database before deleting it.
Keep Audit Log in DB For Last (entries)	The maximum number of audit log entries to be stored in the database. If an entry becomes older than the number of days specified in the Keep Audit Log For field, it is deleted even if the log has fewer than the maximum number of entries.
Purge Now	To delete the old entries from the database, click **Purge Now**. Based on the values entered in the Keep audit log in DB for Last (days) and Keep audit log in DB for Last (entries) fields, the audit log entries that is maximum gets deleted. For example, if the value for **Keep audit log in DB for Last (days)** is 5 and **Keep audit log in DB for Last (entries)** is 5000, and last 5 days log entries is more than 5000, then Cisco Security Manager keeps the last 5000 entries and deletes the older entries even if they belong to the last 5 days. **Note** The Purge Now button only removes audit report entries from the database. It does not remove the *.csv files from the ***<*install_dir*>**\CSCOpx\MDC\log\audit folder. These *.csv files can be deleted directly.
Keep Audit Log File For (days)	Number of days an audit log file should be retained in the system.
Purge Now	Click this button to immediately delete the specified older audit log files. **Note** The Purge Now button deletes the CSDL Audit log files from ***<install_dir>***\MDC\log\CSDL Audit Log folder along with Audit log files from ***<install_dir>***\MDC\log\audit folder.
Keep Operation Log Files For (days)	The number of days that Security Manager keeps operation logs before deleting them. These logs are used for debugging purposes.
Purge Now button	Click this button to immediately delete the specified older operation log files.
Scheduled Purging	
Enable Scheduled Purging for Audit Log Database Entries	Click this check box to scheduled the purging of older log entries. The schedule options are enabled on clicking this check box.
Keep Audit Log in DB For (days)	The maximum number of days of Audit Log entries that has to be stored in the database before deleting it.
Keep Audit Log in DB For Last (entries)	The maximum number of audit log entries to be stored in the database. If an entry becomes older than the number of days specified in the Keep Audit Log For field, it is deleted even if the log has fewer than the maximum number of entries.

Table 11-34 Logs Page (continued)

Enable Scheduled Purging for Audit Log Files	Click this check box to schedule purging of audit logs from the system.
Keep Audit Log File For (days)	Number of days an audit log file should be retained in the system.
Enable Scheduled Purging for Operation Log Files	Click this check box to schedule purging of operation logs from the system.
Keep Operation Log Files For (days)	The number of days that Security Manager keeps operation logs before deleting them. These logs are used for debugging purposes.
Log Level	The level of information, according to severity, that you would like collected in the operation logs. Each level collects different amounts of data. For example, the Info level yields the most data, and the Severe level collects the least.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Policy Management Page

Use the Policy Management page to select the types of router and firewall policies you will manage in Security Manager. These selections apply to routers and firewall devices, but do not apply to IPS devices. By default, all policies are selected for management.

Unmanaged policies are removed from both Device view and Policy view. Any unmanaged policies, local or shared, are removed from the Security Manager database. The only exception is interface policies, which continue to appear in Security Manager but are marked as read-only policies. For firewall devices, interface and failover settings are considered a unit and are managed or unmanaged together.

For detailed information on managing and unmanaging policy types, including what you should do before and after changing these settings, see Customizing Policy Management for Routers and Firewall Devices, page 5-11.

Caution If you use AUS or CNS to deploy configurations to ASA or PIX devices, be aware that the device downloads a full configuration from AUS or CNS. Thus, reducing the policies managed by Security Manager actually removes the configurations from the device. If you intend to deselect some ASA/PIX policies for management to use other applications along with Security Manager to configure devices, do not use AUS or CNS.

Navigation Path

Select **Tools > Security Manager Administration** and select **Policy Management** from the table of contents.

Field Reference

Table 11-35 Policy Management Page

Element	Description
Policies to Manage	The policy types are organized in folders, with router and firewall (which includes all ASA, PIX, and FWSM devices) handled separately, and then by category (NAT, Interfaces, and Platform). Select or deselect policy types as desired and click **Save**. Deselecting the check box for a group of policies deselects all policies in that group. By default, all policies are selected. **Note** Beginning with version 4.18, Cisco Security Manager provides support for ASA 9.10(1) devices that are configured on the Umbrella server.
Display a warning on all shared policies and imported objects	Whether to add a message to all shared policies and to objects that were imported using the **File > Import** command. If you select this option, messages appear on the following: • All shared policies, whether they were imported or locally created. • Policy objects that were created by importing devices or shared policies using the **File > Import** command, but not imported policy objects created by the PolicyObjectImportExport.pl command (described in Importing and Exporting Policy Objects, page 6-23. If you regularly import shared policies, the imported policies and objects replace any same-named policies and objects, so any changes made locally are removed. This message can notify users that policies might be imported and help users identify policy objects that they might not want to edit. **Tip** When importing policies or devices, you are prompted to select a setting for this option. Thus, users who import policies or devices can change this setting without accessing this page provided they have the required authorization. The change is effective only after the importer submits (and if necessary, approves) the changes. For more information, see Importing Policies or Devices, page 10-13.
Save button	Saves your changes. If you are unmanaging a policy, you are shown a list of devices that have the policy assigned to them. Security Manager must be able to obtain the required locks to unassign the policy from all devices, or you must manually unassign the policies (or remove the locks) before unmanaging the policy. If you are managing a previously unmanaged policy, be sure to rediscover all affected devices to bring the existing configurations into Security Manager.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Policy Objects Page

Use the Policy Objects page to define system defaults related to policy object creation.

Navigation Path

Select **Tools > Security Manager Administration** and select **Policy Objects** from the table of contents.

Related Topics

- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Chapter 6, "Managing Policy Objects"

Field Reference

Table 11-36 Policy Objects Page

Element	Description
When Redundant Objects Detected	The action you want Security Manager to take when you try to create a policy object that has the same definition as an existing object: • Ignore—You can freely create objects with identical definitions. Any conflicts are ignored by Security Manager. • Warn—Security Manager displays a warning if you attempt to create an object that is identical to an existing object. You may proceed to create the object, if you wish. • Enforce—Security Manager prevents you from creating an object that is identical to an existing object. An error message is displayed.
Default Source Ports	The port range value that is used as the default source port range for service objects. You can choose one of the following: • Use all ports—Includes all ports from 1 to 65535. • Use secure ports—Includes all ports from 1024 to 65535. If you change the default source ports, you must manually redeploy any previously deployed devices that might be affected. These changes might not be reflected in any open activities until you refresh the data. For more information on port list objects, see Configuring Port List Objects, page 6-102.
Enable AutoComplete Dropdown Box	Whether to have Security Manager list matching service and port list names as you type them when you create a service. You can then easily select from names you have already defined. If you deselect AutoComplete, you have to remember the complete service and port list names and type them in yourself.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Process Monitoring Settings Page

Use the Process Monitoring Settings page to enable process monitoring. Here, you can enable or disable monitoring for specific processes and configure notification settings such as monitoring interval and email addresses. This will send an email notification to specified recipients, when a process stops.

Before You Begin

Configure SMTP Server and sender mail in the CS web console, to get email alerts.

Navigation Path

Select **Tools > Security Manager Administration** and select **Process Monitoring Settings** from the table of contents.

Field Reference

Table 11-37 Process Monitoring Settings Page

Element	Description
Enable Process Monitoring	When selected, Security Manager will allow you to specify the processes that you want to monitor. You must proceed to configure other process monitoring settings. By default, the Process Monitoring feature is disabled on the Cisco Security Manager Server. **Note** Enabling or disabling the process monitor modifies the Windows registry and may generate a system alert.
Monitoring Interval (in minutes)	Specify the interval, at which, the process will be monitored. Valid values are between 1-60 minutes. The default monitoring interval is 5 minutes. **Note** If the monitoring interval is changed, the monitoring task in progress stops and the new monitoring task starts with the updated interval.
Notification Recipient(s) E-mail(s)	Enter email IDs for the notification recipients.You can enter multiple email IDs, separated by a comma. These are the recipients that will be notified, when a process being monitored, stops.
Maximum Mail Alerts	Enter the maximum number of emails that will be sent to the recipients during the course of Security Manager runtime. The default value here, is 10.
Process List	Select one or more processes that you want to monitor. Whenever any of these selected processes stop, a notification email will be sent to the specified recipients. **Note** When Event Management, Health and Performance Monitor and Report Manager are disabled from **Tools > Security Manager Administration,** email notifications will not be sent, even if, the VmsEventServer, CsmHPMServer and CsmReportServer processes are enabled in the process monitoring settings page.
Save button	Saves your changes.

Table 11-37 Process Monitoring Settings Page (continued)

Element	Description
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Single Sign-on Configuration Page

Use the Single Sign-on Configuration page of the Security Manager Administration window to enable and configure a "single sign-on" (SSO) shared key to use for cross-launching Cisco Prime Security Manager or FireSIGHT Management Center.

Note Single sign-on allows users to cross-launch Prime Security Manager or FireSIGHT Management Center from Security Manager without logging into Prime Security Manager or FireSIGHT Management Center separately. However, SSO is not required to cross-launch Prime Security Manager or FireSIGHT Management Center.

Tip Cisco Prime Security Manager is used to manage ASA CX modules. FireSIGHT Management Center is used to manager ASA FirePOWER modules.

Related Topics

- Detecting ASA CX and FirePOWER Modules, page 72-21
- Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20
- Sharing Device Inventory and Policy Objects with PRSM, page 72-22

Navigation Path

1. Click **Tools > Security Manager Administration** and select **Single Sign-on Configuration** from the table of contents.
2. Select **Enable for Prime Security Manager** [checkbox] or **Enable for FireSIGHT Management Center** [checkbox].

Field Reference

Table 11-38 Single Sign-on Configuration Page

Element	Description
Enable for Prime Security Manager	[checkbox] Lets you enable or disable the SSO feature for Prime Security Manager. When disabled, the shared key is retained.
Enable for FireSIGHT Management Center	[checkbox] Lets you enable or disable the SSO feature for FireSIGHT Management Center. When disabled, the shared key is retained.

Table 11-38 Single Sign-on Configuration Page (continued)

Element	Description
Shared Key for Single Sign-on	Use the features in this section to generate and view an encryption key for cross-launching Prime Security Manager or FireSIGHT Management Center. Click the **Generate** button to randomly generate a 128-bit AES key, which is then displayed as a 32 hexadecimal string in the SSO Shared Key field. **Note** This key must be provided when configuring single sign-on cross-launching in Prime Security Manager or FireSIGHT Management Center. Also, each allowed Security Manager user must be configured in the Prime Security Manager database or the FireSIGHT Management Center user database with the same username as that in the Security Manager user database (the password can be different). **Tip** Refer to "Configuring Single Sign-On for Cisco Security Manager" in the *User Guide for ASA CX and Cisco Prime Security Manager* (Cisco ASA CX Context-Aware Security End-User Guides) for information about configuring SSO in PRSM.

Rule Expiration Page

Use the Rule Expiration page to define the default values for policy rule expiration. When you create policies for some types of policy rules (such as access rules), you can set an expiration date for the rule, and Security Manager can notify you by e-mail of the approaching expiration date.

You must configure an SMTP server to enable e-mail notifications. For more information, see Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.

Navigation Path

Select **Tools > Security Manager Administration** and select **Rule Expiration** from the table of contents.

Field Reference

Table 11-39 Rule Expiration Page

Element	Description
Notify Email	The default e-mail address that should receive notifications of rule expiration. Users can override this address when configuring individual rules.
Notify Before Expiration	The default number of days before a rule expires that Security Manager should send the e-mail message. Users can override this value when configuring individual rules.
Sender	The e-mail address that Security Manager will use for sending e-mail notifications.

Table 11-39 Rule Expiration Page (continued)

Element	Description
Email Format	The format of the e-mail message: • Text—The e-mail is sent in HTML and plain text formats. • XML—The e-mail is sent using an XML markup. This option might be appropriate if you decide to write a program to automatically process and respond to notifications.
Save button	Saves your changes.
Reset button	Restores all fields to their previous values.
Restore Defaults button	Resets values to Security Manager defaults.

Server Security Page

Use the Server Security page to open specific pages in the CiscoWorks Common Services application, where you can configure various security features on the Security Manager server. CiscoWorks Common Services controls the basic functions of the Security Manager server, including user access control and system security.

When you log in to Security Manager, your username and password are compared with the account information stored in the CiscoWorks or Cisco Secure Access Control Server (ACS) database, depending on which system you established at installation as your AAA provider. After the authentication of your credentials, you have access according to the role you have been assigned.

For more information on Security Manager roles and privileges, including descriptions of how Common Services roles translate to user functions in Security Manager, see the *Installation Guide for Cisco Security Manager*.

Navigation Path

Select **Tools > Security Manager Administration** and select **Server Security** from the table of contents.

Field Reference

Table 11-40 Server Security Page

Element	Description
AAA Setup button	Opens Common Services and displays the AAA Mode Setup page. From this page, you can set AAA as your fallback sign-on method. For more information about AAA, click **Help** from the AAA Mode Setup page.
Certificate Setup button	Opens Common Services and displays the Self-Signed Certificate Setup page. CiscoWorks enables you to create self-signed security certificates, which you can use to enable SSL connections between your client browser and management server. For more information about self-signed certificates, click **Help** from the Certificate Setup page.

Table 11-40 Server Security Page (continued)

Element	Description
Single Sign On button	Opens Common Services and displays the Single Sign-On Setup page. With Single Sign On (SSO), you can use your browser session to transparently navigate to multiple CiscoWorks servers without having to authenticate to each of them. Communication between multiple CiscoWorks servers is enabled by a trust mode addressed by certificates and shared secrets. For more information about setting up SSO, click **Help** from the Single Sign-On page.
Local User Setup	Opens Common Services and displays the Local User Setup page, from which you can add and delete users, edit user settings, and assign roles or permissions. For more information, click **Help** from the Local User Setup page and see the *Installation Guide for Cisco Security Manager*.
System Identity Setup	Opens Common Services and displays the System Identity Setup page. Communication between multiple CiscoWorks servers is enabled by a trust mode addressed by certificates and shared secrets. System Identity setup helps you to create a trust user on servers that are part of a multi-server setup. For more information about system identity setup, click **Help** from the System Identity Setup page.
Native RBAC Parameters	
Allow logon for user ids not available in Local User Database	For Security Manager installations integrated with an external authentication server like Active Directory, TACACS+, or RADIUS, specifies whether users can log in even when their user name is not defined in the Security Manager user list. When enabled, users are allowed to log in using the default role specified in Role Management Setup. If a default role is not configured, the user is not allowed to log in.

Take Over User Session Page

Use the Take Over User Session page to take over another user's configuration session. A user with administrative privileges can take over the work of another user in non-Workflow mode. Taking over a session is useful when a user is working on devices and policies, causing the devices and policies to be locked, and another user needs access to the same devices and policies. However, when you take over another user's session, your current session is discarded, so make sure that you submit your changes before taking over a session.

The table shows all current configuration sessions, listing the user name and the state of the session, whether the user is currently logged in or logged out. Select the configuration session you want to take over and click **Take over session**. The session is transferred to you in its current state, including any saved changes the user made during the session.

If the selected user is logged in at the time you take over the session, the user receives a warning message, loses any unsaved changes in progress, and then is logged out.

For more information, see Taking Over Another User's Work, page 10-23.

Navigation Path

Select **Tools > Security Manager Administration** and select **Take Over User Session** from the table of contents.

Ticket Management Page

Use the Ticket Management page to enable Ticket Management, to configure a ticketing system URL for integration with an external change management system, and to configure purge settings for ticket information.

When Ticket Management is enabled, every Image Management installation job must have an assigned ticket or it will not be performed.

Navigation Path

Select **Tools > Security Manager Administration** and select **Ticket Management** from the table of contents.

Related Topics

- Changing Workflow Modes, page 1-28
- Comparing Workflow Modes, page 1-22

Field Reference

Table 11-41 Ticket Management Page

Element	Description
Enable Ticketing	Whether to enable Ticket Management.
System Generated Default Ticket Name	By default, this check box is checked. Clear the check box, if you do not want the ticket name to be appended with the system generated default name. The ticket name field in the activity creation dialog is left blank.
Ticketing System URL	
Ticketing System URL	The URL to use for launching an external change management system. When this field is configured, the Ticket ID is a hyperlink that will launch the URL specified. The URL must be formatted as a template that accepts the Ticket ID as part of the URL. The template format uses {0} in place of the actual Ticket ID. For example, if the URL to launch an external ticket management system for a ticket with the ticket ID of *TKT12345* is http://ticketsystem/displayticket?ticketid=TKT12345, then the template URL you would use would be **http://ticketsystem/displayticket?ticketid={0}**. When you create a ticket, the Ticket ID you specify will be used in the hyperlink in place of the {0}.
Generate	Click to display the Generate Template URL dialog box that can be used to create a Ticketing System URL. Using the example above, you would enter **TKT12345** in the Ticket ID field and **http://ticketsystem/displayticket?ticketid=TKT12345** in the Ticket URL field. When you click **OK**, the appropriate template URL is created and entered into the Ticketing System URL field.

Table 11-41 Ticket Management Page (continued)

Element	Description
Ticket History	
Ticket History settings are only available in non-Workflow mode. In Workflow mode, purge settings are controlled via the settings for Activities (see Workflow Page, page 11-75).	
Purge Tickets (including change report) Older than	The number of days that ticket information should be kept in the Ticket Manager table. The default is 30. You can specify from 1 to 120 days. Click **Purge Now** to delete all tickets older than the number of days specified.
Purge Change Report older than	The number of days that change reports should be maintained. The default is 30. You can specify a value that is less than the Purge Tickets (including change report) Older than setting. Click **Purge Now** to delete all change reports older than the number of days specified.
Save button	Saves your changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

Token Management Page

Use the Token Management page to identify the Token Management System (TMS) server to use for deploying configurations to Cisco IOS routers that use TMS as the communication protocol. Security Manager uses the settings on this page to contact the TMS server.

Security Manager uses FTP to deploy the delta configuration file to the TMS server, from which the configuration file can be downloaded and encrypted onto an eToken.

To use TMS with Cisco IOS routers, you must specify TMS as the transport protocol. You can do this for all routers on the Device Communication page (see Device Communication Page, page 11-21), or for a specific router in its device properties (see Device Properties: General Page, page 3-41). You must also configure the TMS server as an FTP server, otherwise deployment will fail.

Navigation Path

Select **Tools > Security Manager Administration** and select **Token Management** from the table of contents.

Related Topics

- Deploying Configurations to a Token Management Server, page 8-43
- Understanding Deployment Methods, page 8-8

Field Reference

Table 11-42 Token Management Page

Element	Description
Server Name or IP Address	The DNS hostname or IP address of the TMS server.

Table 11-42 Token Management Page (continued)

Element	Description
Username	The username Security Manager should use to log on to the TMS server.
Password Confirm Password	The password for the username. Enter the password in both fields.
Directory in the TMS Server for Config Files	The directory on the TMS server where deployed configuration files will be downloaded. The root FTP directory (".") is the default FTP location on the TMS server.
Public Key File Location	The location of the public and private key files on the Security Manager server, as copied from the TMS server. Security Manager uses the public key to encrypt data sent to the TMS server. Then the server uses its private key to decrypt the data. Security Manager comes with a default public key that matches the default private key on the server. **Note** If needed, you can generate a new pair of public and private keys using the TMS server. If you do this, you need to copy the new public key to the Security Manager server.
Save button	Saves your changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

VPN Policy Defaults Page

Use the VPN Policy Defaults page to view or assign the default VPN policies that Security Manager uses for each IPsec technology. Before you can select a policy as a default, you must create the policy as a shared policy, submit it to the database and have it approved. You cannot create policies from this page. For detailed information on how to configure these defaults, see Understanding and Configuring VPN Default Policies, page 25-12.

For each tab that relates to a VPN topology, the drop-down lists for each policy type list the existing shared policies that you can select. You can select a policy and click the **View Content** button to see the definition of that policy. In some cases, you are allowed to make changes, but you cannot save them.

Security Manager uses VPN policy defaults to simplify VPN configuration while ensuring that policy consistency is maintained. Security Manager provides factory default policies for mandatory policies, which provide values for settings that must be configured on the devices in your VPN topology for the VPN to work. Mandatory policies differ depending on the assigned IPsec technology. Factory default policies with their default configurations enable you to deploy to your devices immediately after creating the VPN topology. Default settings are not provided for optional policies. You might want to create shared policies to provide different default settings instead of using the factory default settings.

Navigation Path

Select **Tools > Security Manager Administration** and select **VPN Policy Defaults** from the table of contents.

Related Topics

- Assigning Initial Policies (Defaults) to a New VPN Topology, page 25-62

- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25
- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices), page 30-36

Field Reference

Table 11-43 VPN Policy Defaults Page

Element	Description
DMVPN tab	Lists the policy types for which you can configure default policies for the Dynamic Multipoint VPN technology.
Large Scale DMVPN tab	Lists the policy types for which you can configure default policies for the Large Scale Dynamic Multipoint VPN technology.
Easy VPN tab	Lists the policy types for which you can configure default policies for the Easy VPN technology.
IPsec/GRE tab	Lists the policy types for which you can configure default policies for the IPsec/GRE VPN technology.
GRE Dynamic IP tab	Lists the policy types for which you can configure default policies for the GRE Dynamic IP VPN technology.
Regular IPsec tab	Lists the policy types for which you can configure default policies for regular IPsec VPN technology.
Regular IPsec VTI tab	Lists the policy types for which you can configure default policies for regular, tunnel-based IPSEC VPN technology.
GET VPN	Lists the policy types for which you can configure default policies for the Group Encrypted Transport (GET) VPN technology.
Remote Access VPN	Lists the policy types for which you can configure default policies for IPsec remote access VPNs.
S2S Endpoints tab	The interface roles that define the default endpoints for internal and external interfaces in site-to-site VPNs.

Workflow Page

Use the Workflow page to select the workflow mode that Security Manager enforces and to define the default settings for activity and deployment job notifications and logging.

Before changing the workflow mode, read the following topics to understand how the modes differ and the effects of changing the modes:

- Working in Workflow Mode, page 1-21
- Working in Non-Workflow Mode, page 1-22
- Comparing Workflow Modes, page 1-22
- Changing Workflow Modes, page 1-28

Navigation Path

Click **Tools > Security Manager Administration** and select **Workflow** from the table of contents.

Related Topics

- Chapter 4, "Managing Activities"
- Chapter 8, "Managing Deployment"

Field Reference

Table 11-44 Workflow Page

Element	Description
Workflow Control	
Enable Workflow	Whether to enable Workflow mode. When Workflow mode is enabled, you can select whether or not to have an approver for activities and deployment jobs.
Require Activity Approval	Whether to require that activities be approved explicitly by an assigned approver. For more information about the differences between working with and without an approver, see Activity Approval, page 4-3.
Submitter can Approve Activity	Activities can be approved by submitter.
Require Deployment & Install Image Approval	Whether to require that deployment jobs and install image jobs be approved explicitly by an assigned approver. For more information about the differences between working with and without an approver, see Understanding Deployment, page 8-1.
Submitter can Approve Deployment Jobs	Deployment jobs can be approved by submitter.
System Generated Default Activity Name	By default, this check box is checked. Clear the check box, if you do not want the activity name to be appended with the system generated default name. The activity name field in the activity creation dialog is left blank.
Email Notifications	
Sender	The e-mail address that Security Manager will use for sending e-mail notifications.
Activity Approver	The default e-mail address for the person responsible for approving activities. Users can override this address when submitting an activity for approval. For more information, see Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20.
Job/Schedule Approver	The default e-mail address of the person responsible for approving deployment jobs or schedules. Users can override this address when submitting a job or schedule for approval. For more information, see Submitting Deployment Jobs, page 8-38.
Require Deployment Status Notification Include Job Deployer Job Completion Notification	Whether to have e-mail notifications sent whenever the status of a deployment job changes. If you select this option, enter the e-mail addresses that should receive notification in the Job Completion Notification field. Separate multiple addresses with commas. You can also select Include Job Deployer to include the e-mail address of the person who deployed the job on the notification e-mail message.

Table 11-44 Workflow Page (continued)

Element	Description
Workflow History	
Keep Activity for	The number of days that activity information should be kept in the Activity table. The default is 30. You can specify from 1 to 180 days. Click **Purge Now** to delete all activities older than the number of days specified. **Note** If ticketing is enabled in non-Workflow mode, purge settings are controlled via the settings for Tickets (see Ticket Management Page, page 11-72).
Keep Job for	The number of days that deployment job information should be kept in the Deployment Job table. The default is 30. You can specify from 1 to 180 days. Click **Purge Now** to delete all jobs older than the number of days specified.
Keep job per schedule for	The number of days that deployment job information should be kept in the Deployment Job table for each job schedule. This setting applies only to jobs that were initiated by a schedule. The default is 30. You can specify from 1 to 180 days. Click **Purge Now** to delete all jobs older than the number of days specified.
Save button	Saves your changes.
Reset button	Resets changes to the previously applied values.
Restore Defaults button	Resets values to Security Manager defaults.

Wall Settings Page

The Security Manager Wall Settings page is where you can enable or disable the Wall feature.

The "Wall" feature is also called the "ShoutBox" feature. You can use it to send messages to all users who are logged in on the same Security Manager server. First, however, it must be enabled on the Wall Settings page.

Note Only admin users have permission to enable or disable the Wall feature, but all users have permission to send messages.

You would want to use the Wall feature, for example, to interact with other users while making some changes in your Security Manager installation, perhaps about the changes being made or certain immediate actions to be performed on the changes. The message being sent is broadcast to all users who are logged in. The Wall feature allows users to enter basic profile information that can be viewed by others when logged in. A significant use of the Wall feature is that it can be used to view a list of all users who are currently logged in. (A user is removed from the Wall window after idle timeout or logging out through the Security Manager client.)

You cannot use the Wall feature to send *.pdf, *.xls, or other file attachments.

Navigation Path

Click **Tools > Security Manager Administration** and select **Wall Settings** from the table of contents.

Field Reference

Table 11-45 Wall Settings Page

Element	Description
Enable users to send messages to others	Whether to enable or disable the Wall feature.
Save button	Saves and applies changes.
Reset button	Resets changes to the last saved values.
Restore Defaults button	Resets values to Security Manager defaults.

When the Wall feature is enabled, you can open the Wall window by clicking **Tools > Wall...** or by clicking the Wall icon in Configuration Manager.

You can also open the Wall window by clicking the Wall icon in Health and Performance Monitor or Image Manager. You cannot open the Wall window in Event Viewer or Report Manager.

Detailed Wall feature help is available on the Wall window by clicking the help icon.

The Wall window contains the following elements:

- Left-hand pane, which shows the users who are logged in on the same Security Manager server and an expand/collapse button.
- Right-hand pane, which occupies most of the page and contains the text of the messages that users have sent. The right-hand pane also has a button to enable or disable wall alerts and the help icon, which you can click to see detailed help.

Summary of the Wall Window	
Message Display	Messages are displayed in the right-hand pane of the Wall window. The latest messages are always displayed at the top. Selective copy of text is allowed from the messages. You are allowed to type a maximum of 280 characters in the message panel, and after completing this number of characters you are alerted by a beep sound.
Message Log	You can see a log of previous messages. The message log keeps 100 messages. The messages become visible to you when you launch the Wall window.

Profile Picture	You can upload a picture for your profile. Valid image types such as JPG, PNG, BMP, and GIF types are supported. To upload a picture, use the **upload** link in the user profile window. To open the user profile window, click on the username or user picture in the Wall window. The user profile window also has an icon that toggles between Edit profile information and Save profile information.
User Profile Window	To open the user profile window, click on the username or user picture in the Wall window. The user profile window contains the following information: • Profile Name (maximum 20 characters) • Designation (maximum 15 characters) • Email (maximum 15 characters) • Phone Click the corresponding email link to send the mail.
Notification Alert	On receipt of a new message and when the Wall window is not focused, a new notification alert popup is shown to you. You can simply click on the notification to launch the Wall window. When the notification alert popup is shown to you, the Wall window icon also flashes with the message count displayed. You can turn off the notifications alert from the settings options provided on the alert popup or in the Wall window.

Figure 11-1 Wall Window

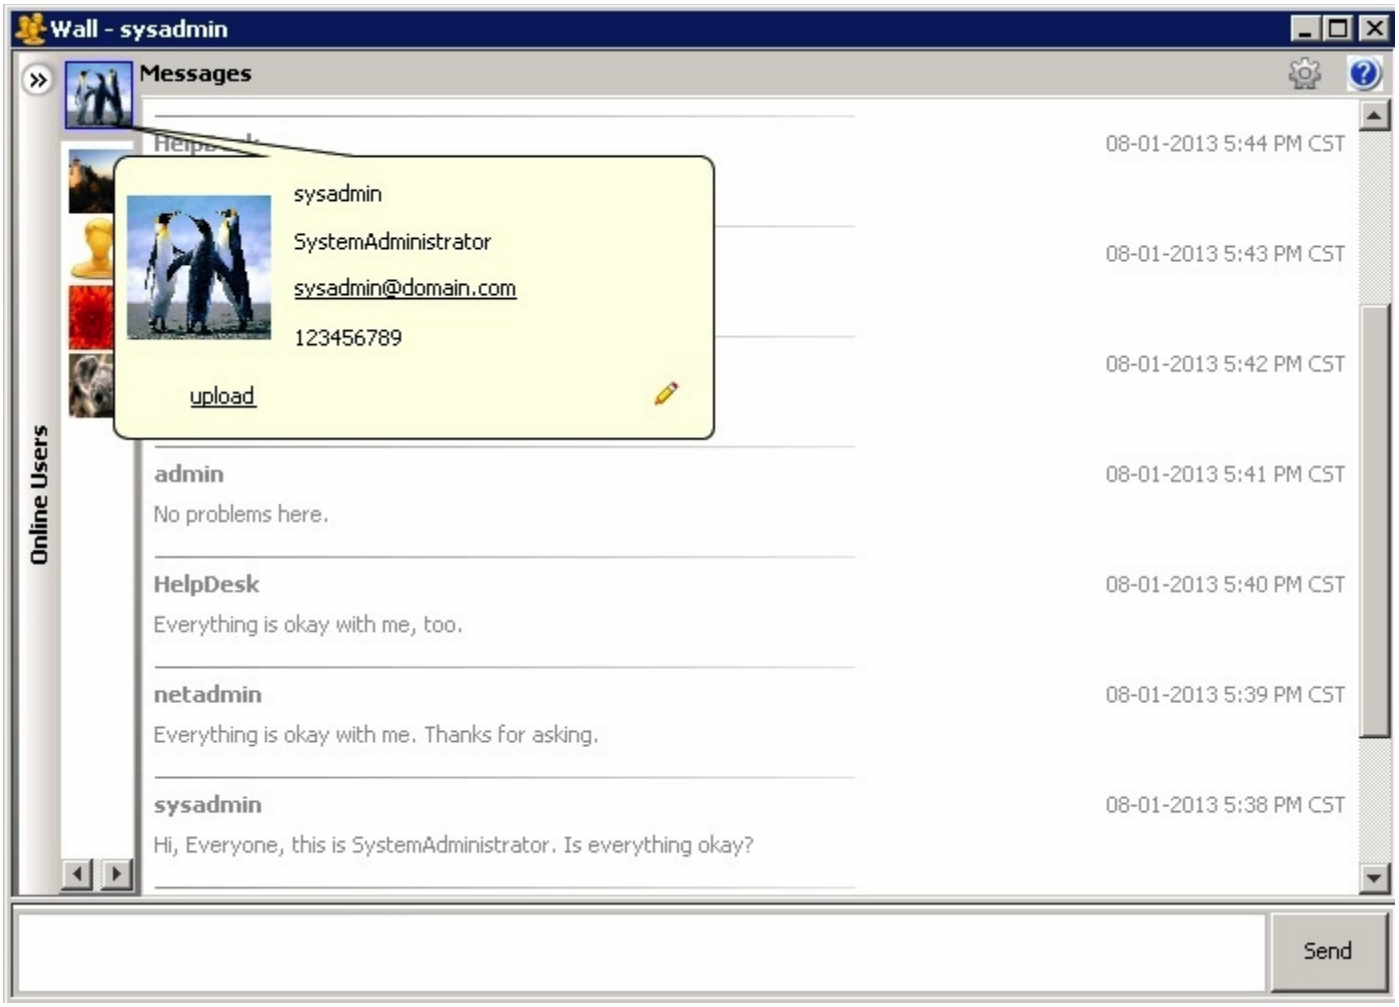

PART 2

Firewall Services and NAT

Introduction to Firewall Services

The Firewall policy folder (in either Device or Policy view) includes firewall-related policies that you can deploy to the Adaptive Security Appliance (ASA), PIX Firewall (PIX), Catalyst Firewall Services Module (FWSM), and security routers running Cisco IOS Software. These policies allow you to control network access through a device.

This chapter contains the following topics:

- Overview of Firewall Services, page 12-1
- Managing Your Rules Tables, page 12-7

Overview of Firewall Services

The Firewall policy folder (in either Device or Policy view) includes firewall-related policies that you can deploy to the Adaptive Security Appliance (ASA), PIX Firewall (PIX), Catalyst Firewall Services Module (FWSM), and security routers running Cisco IOS Software, including Aggregation Services Routers (ASR) and Integrated Services Routers (ISR).

These policies are focused on controlling access through the device, rather than access to the device (that is, logging into the device so that you can change its configuration or use **show** commands). Following is a general overview of the available firewall policies with pointers to topics that provide more detailed information:

- AAA rules—These are AAA firewall or authentication proxy rules that can require a user to authenticate (with a username and password) and optionally be authorized before the device allows the user to make network connections through it. You can also create accounting rules to collect billing, security, or resource allocation information. For more information, see Understanding AAA Rules, page 15-1.
- Access rules—These are traditional interface-based extended access control rules. They permit or deny a packet based on source address, destination address, source interface, and service, and you can apply them in both the in and out directions. For more information, see Understanding Access Rules, page 16-1.
- Inspection rules—These are traditional Context-Based Access Control (CBAC) inspection rules that filter out bad TCP/UDP packets based on application-layer protocol session information and that enable return traffic for the selected services. For more information, see Understanding Inspection Rules, page 17-2.
- Web filter rules—These are a type of inspection rule that filters web traffic based on the requested URL, allowing you to prevent connections to undesirable web sites. For more information, see Understanding Web Filter Rules, page 18-1.

- Zone-based firewall rules—These rules replace access rules, inspection rules, and web filter rules on IOS devices if you want to configure your rules based on zones instead of interfaces. A zone is a defined group of interfaces that perform the same security role (such as Inside or Outside). By using zone rules, you can create more compact device configurations than you can by using the other types of rules. For more information, see Understanding the Zone-based Firewall Rules, page 21-3.
- Botnet Traffic Filter Rules—These rules help you to spot botnet traffic when it is sent to known bad addresses. Botnets install malware on unsuspecting computers and use those computers as proxies to perform malicious actions. For more information, see Chapter 19, "Managing Firewall Botnet Traffic Filter Rules".
- Transparent rules—These are Ethertype access control rules that apply to non-IP layer-2 traffic on transparent or bridged interfaces. For more information, see Configuring Transparent Firewall Rules, page 23-1.

Most firewall rules policies are configured in rules tables. These tables allow in-line editing for most cells, rule organization using sections, and the ability to change the order of rules. If you create shared rules policies, you can apply them to a number of devices, even to devices running different operating systems, and Security Manager automatically creates the appropriate device commands to configure the policies based on the characteristics of each individual device, filtering out settings that do not apply to a device. For more information on using rule tables, see Managing Your Rules Tables, page 12-7.

Another powerful feature used by most firewall rules policies is the idea of inheritance. When you create shared policies, one of your options is to have a device inherit the policy rather than be assigned the policy. This allows you to have a set of shared rules that apply to all devices, while having unique rules that apply to only those devices that require them. For more information about inheritance, see the following topics:

- Understanding Rule Inheritance, page 5-4
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37.

The following topics provide additional overview information about firewall services policies:

- Understanding the Processing Order of Firewall Rules, page 12-2
- Understanding How NAT Affects Firewall Rules, page 12-3
- ACL Names Preserved by Security Manager, page 12-4

Understanding the Processing Order of Firewall Rules

When you configure firewall rules policies, you should keep in mind the logical order in which the rules are processed. For example, if you plan to drop all traffic of a certain type in an access rule, there is no reason to create rules in other firewall policies that apply to that type of traffic, because they will never be triggered. Conversely, if you want to apply certain types of inspection or web filtering on traffic, you must ensure that your access rules first allow that traffic to enter the device.

Following is the general logical processing order of firewall rules:

- AAA rules—If you require authentication, with or without authorization, the user must successfully pass the test or the traffic is dropped.
- Access rules (In direction)—The traffic must then get through your access rules. If you used AAA rules, you might have allowed temporary per-user access rules to be configured for the user's session. These per-user rules are configured in your AAA server, not in Security Manager.

  On ASA 8.3+ devices, global access rules are then processed after any interface-specific access rules. For more information, see Understanding Global Access Rules, page 16-3.

- One of the following:
  - Inspection rules (In direction), web filter rules (In direction), botnet rules, service policy rules (IPS, QoS, Connection)—All of these are applied to the traffic. For devices that do not allow you to configure the direction, all rules are considered to be in the In direction.
  - Zone-based firewall rules—If you configured zone-based rules for an IOS device, these rules replace inspection and web filter rules (botnet rules do not apply to IOS devices).
- Routing protocols are then applied to the traffic. The traffic is dropped if it cannot be routed. (Routing policies are in the Platform folders for the various device types and are not considered firewall policies.)
- ScanSafe Web Security policies, Inspection rules (Out direction), web filter rules (Out direction)—For IOS devices only, if you created ScanSafe policies, or inspection or web filter rules in the Out direction, they are now applied.
- Access rules (Out direction)—Finally, the traffic must pass through any Out direction access rules.

Transparent rules do not fit into this picture. Because transparent rules apply to non-IP layer-2 traffic only, if a transparent rule applies to a packet, no other firewall rule applies to it; and conversely, if other rules apply, the transparent rule never applies.

Related Topics

- Understanding AAA Rules, page 15-1
- Understanding Access Rules, page 16-1
- Understanding Inspection Rules, page 17-2
- Understanding Web Filter Rules, page 18-1
- Understanding the Zone-based Firewall Rules, page 21-3
- Chapter 19, “Managing Firewall Botnet Traffic Filter Rules”
- Configuring Transparent Firewall Rules, page 23-1

Understanding How NAT Affects Firewall Rules

Devices that support firewall rules also allow you to configure network address translation (NAT). NAT substitutes the real address in a packet with a mapped address that is routable on the destination network.

If you configure NAT to occur on an interface, the firewall rules that are also configured on that interface should assess traffic based on the translated address rather than on the original (pre-NAT) address, with the exception of ASA 8.3+ devices.

Devices running ASA software release 8.3 and higher use the original, or real, IP address when evaluating traffic with the exception of IPSec VPN traffic policies. Thus, when you configure firewall rules, ACL policy objects, or the IOS, QoS, and connection rules platform service policy, ensure that you use the original addresses.

For more information about NAT, see the following topics:

- ASA, PIX, FWSM devices—Understanding Network Address Translation, page 24-2.
- IOS devices—NAT Policies on Cisco IOS Routers, page 24-5.

ACL Names Preserved by Security Manager

Security Manager tries to preserve user-defined access control list (ACL) names as they appear in device configurations. Security Manager can preserve the ACL names configured on a device in the following circumstances:

- If the ACL name is specified in Security Manager.

  For access rules policies, you can specify ACL names in **Firewall > Settings > Access Control** or **Firewall > Settings > IPv6 Access Control**. You can specify a given name for a single interface and direction, but the name is used for any other interfaces and directions that use the same ACL. Keep in mind that you cannot use the same name as an ACL policy object that you assign to other policies on the device, and you cannot use the same name for IPv4 and IPv6 ACLs.

  Note Prior to the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, separate pages, policies and policy objects were provided for configuring IPv4 and IPv6 firewall rules and policies. With Security Manager 4.4 and ASA 9.0+, these policies and policy objects were combined or unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided in Device view, while in Policy view, IPv4 and unified versions of the AAA-, access- and inspection-rule policy types are provided.

- If a policy uses an ACL policy object, the name of the policy object is used for the ACL name. ACL policy objects created during discovery use the name of the ACL defined on the device whenever possible. Behavior depends on an administrative setting:
  - If you select **Allow Device Override for Policy Objects** in **Tools > Security Manager Administration > Discovery**, if a policy object with the same name exists in Security Manager, but it has different content, the name is reused and a device-level override is created.
  - If you do not select that option, a new policy object is created with the same name but with a number appended to it, for example, ACLobject_1. This is the default behavior.
- If you select **Reuse Existing Names** for the **Firewall Access List Names** setting in **Tools > Security Manager Administration > Deployment**, names defined on the device are reused for firewall rules that generate ACLs.
- If the ACL is unshared, even if you change the content of the ACL in Security Manager.
- If the ACL is shared, but the policies that share the ACL are defined identically in Security Manager. If you change the content of the ACL, one ACL retains the name and the others are assigned generated names.

Note On ASA devices and on PIX devices not running version 6.3(x), Security Manager does not reuse the ACL name if it is used by a NAT policy static rule and contains an object-group. The ACL is deployed with the contents of the object-group defined as the source. This is because the device requires that all ACEs in the ACL have the same source.

Tips

- If you use an ACL policy object that uses a name also used by an ACL already defined on the device, and the existing ACL is for a command that Security Manager does not support, you will get a deployment error asking you to choose a different name. If this happens, rename the policy object.

- ACLs named <number>_<number> are not valid on IOS devices. Security Manager strips off the suffix prior to deployment. This also means that you cannot assign an IOS device more than one ACL object with the same numbered prefix. However, named ACLs that have a numbered suffix are allowed, for example, ACLname_1.
- Numbered ACLs must use the correct number ranges for IOS devices. Standard ACLs must be in the range 1-99 or 1300-1999. Extended ACLs must be in the range 100-199 or 2000-2699.
- ACL names for IOS devices cannot begin with an underscore (_).
- Policies that do not preserve user-defined names include SSL VPN policies, transparent firewall rules, and AAA rules (for IOS devices).

The following topics provide additional information about ACL naming:

- ACL Naming Conventions, page 12-5
- Resolving User Defined ACL Policy Naming Conflicts, page 12-6
- Resolving ACL Name Conflicts Between Policies, page 12-7

ACL Naming Conventions

When the name for the ACL is generated by Security Manager, the name is derived from the type of rule or platform being defined and certain configuration settings that make it unique. All newly created ACLs are given a name based on the naming conventions shown in the following table.

Tip During deployment, sometimes a suffix *.n* (where n is an integer) might get added to an ACL name if the existing ACL cannot be edited in place. For example, if an ACL named acl_mdc_outside_10 already exists on the device, a new ACL with the name acl_mdc_outside_10.1 is created if you do not remove the old ACL before you deploy the new ACL.

Table 12-1 ACL Naming Conventions

Policy Type	Naming Convention
Access ACLs	• Inbound: CSM_FW_ACL_InterfaceName • Outbound: CSM_FW_ACL_OUT_InterfaceName
IPv6 Access ACLs	• Inbound: CSM_IPV6_FW_ACL_InterfaceName • Outbound: CSM_IPV6_FW_ACL_OUT_InterfaceName **Note** Prior to the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, separate pages, policies and policy objects were provided for configuring IPv4 and IPv6 firewall rules and policies. With Security Manager 4.4 and ASA 9.0+, these policies and policy objects were combined or unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided in Device view, while in Policy view, IPv4 and unified versions of the AAA-, access- and inspection-rule policy types are provided.
Inspection Rules	• For ASA 7.0+/PIX 7.0+: CSM_CMAP_ACL_n where n is an integer beginning with 1. • For IOS devices, a numbered ACL.

Table 12-1 ACL Naming Conventions (continued)

Policy Type	Naming Convention
NAT0 ACLs	• Inbound: CSM_nat0_InterfaceName_in • Outbound: CSM_nat0_InterfaceName
NAT ACLs	• Inbound: CSM_nat_InterfaceName_poolID_in • Outbound: CSM_nat_InterfaceName_poolID **Note** For PIX 6.3(x) devices, the following is added to the ACL name: add _dns for dns, _nrseq for norandomseq, _emb## for embryonic limit and _tcp## and _udp## for tcp and udp max connection limits.
NAT Policy Static Translation Rules ACLs	• For PIX 6.3(x) devices: – For IP: CSM_static_globalIP_LocalInterfaceName_globalInterfaceName – For other protocols: CSM_static_globalIP_LocalInterfaceName_globalInterfaceName_ protocol _globalPort • For devices running other OS versions, the localIP string is added: – For IP: CSM_static_localIP_globalIP_LocalInterfaceName_globalInterfaceName – For other protocols: CSM_static_localIP_globalIP_LocalInterfaceName_globalInterfaceName_ protocol _globalPort
AAA ACLs	For PIX/ASA/FWSM: CSM_AAA_{AUTHO | ATHEN | ACCT}_InterfaceName _ServerGroupName Authentication Proxy for IOS devices: • On an interface without NAC: CSM_AUTH-PROXY_InterfaceName_traffic type_ACL, where InterfaceName is the interface in which the rule is applied and traffic type is HTTP, Telnet, or FTP. • AuthProxy and NAC on the same interface: CSM_ADMISSION_ID_ACL, where ID is an internal identifier of the interface role within Security Manager to which NAC is applied.
Web Filter Rules ACLs	For ASA 7.0+/PIX 7.0+: devices correspond to a filter command. For IOS devices, a numbered ACL.

Resolving User Defined ACL Policy Naming Conflicts

Cisco Security Manager generates ACL names that begin with “CSM_”. You should not use the same naming pattern while defining a ACL in the device. If you declare ACL names with the “CSM_” prefix on device, during discovering the device configuration in Cisco Security Manager, those ACL names are replaced with Security Manager generated names and respective configuration delta would be applied to a device on the next deployment.

For example, Cisco Security Manager has CSM_FW_ACL_InterfaceName as the ACL naming pattern for inbound firewall interface. Here, if you use the CSM pattern for ACL name declaration in device like, CSM_xyz, Security Manager renames it as “CSM_FW_ACL_InterfaceName”.

Note This rule is valid for firewall access list (Table 12-1) and delta would be generated and applied to a device even when the *Reuse existing names* setting is configured in Tools > Security Manager Administration > Deployment.

Resolving ACL Name Conflicts Between Policies

If an ACL is shared, but the policies that share the ACL are not defined identically in Security Manager, one policy uses the original name of the ACL and the other policies use a new name generated by Security Manager. The order of preference for determining which policy uses the original name is as follows:

- Access list ACLs
- AAA ACLs
- Static ACLs
- NAT0 ACLs
- NAT ACLs

For example, if an access ACL and a NAT0 ACL try to reuse the same ACL, the access ACL uses the original name as configured on the device and the NAT0 ACL is renamed by Security Manager.

Managing Your Rules Tables

The following sections explain some of the basics of using rules tables, which appear in many of the firewall rules, NAT, and select other policies:

- Using Rules Tables, page 12-8
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Finding and Replacing Items in Rules Tables, page 12-16
- Moving Rules and the Importance of Rule Order, page 12-19
- Enabling and Disabling Rules, page 12-20
- Using Sections to Organize Rules Tables, page 12-20
- Combining Rules, page 12-22
- Generating Policy Query Reports, page 12-28
- Optimizing Network Object Groups When Deploying Firewall Rules, page 12-35
- Expanding Object Groups During Discovery, page 12-35

Using Rules Tables

Rules tables in Security Manager display sets of rules (for example, access rules) that make up a policy. These types of tables are used in only a select group of policies, but many of the firewall services rules policies use them. Rules tables are used when the order of the rules within the policy matter.

Figure 12-1 details the features in rules tables.

Figure 12-1 Rules Table Example

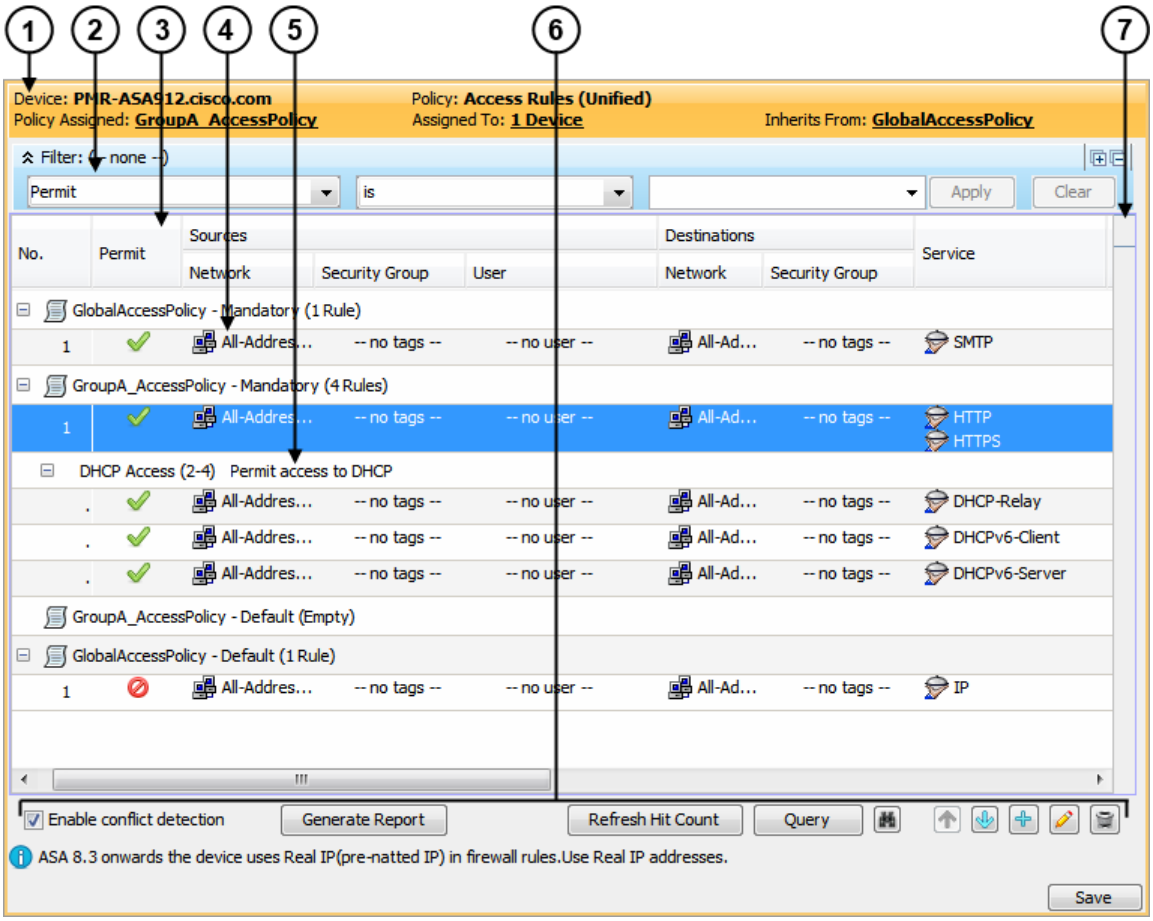

Following is an explanation of the numbered call-outs of the rules table features:

- **Device and policy identification banner (1)**—The banner provides information about policy sharing and inheritance and includes the ability to perform some actions. For detailed information, see Using the Policy Banner, page 5-38.
- **Table filter (2)**—You can filter the rules displayed to help you find rules in a large table. For more information, see Filtering Tables, page 1-48.
- **Table column headings (3)**—You can sort by column and move, show, and hide columns. For more information, see Table Columns and Column Heading Features, page 1-49.
- **Rules, Workarea (4)**—The body of the table shows the rules that are included in the policy.
- **User-defined sections (5)**—You can group rules into sections for your convenience. For more information, see Using Sections to Organize Rules Tables, page 12-20.

- **Table buttons (6)**—Use the buttons below the table to do the following:
  - Enable automatic conflict detection (Access Rules only). For more information, see Using Automatic Conflict Detection, page 16-28.

    If conflict detection is enabled, you can click the **Generate Report** button to create an HTML report of the conflicts that can be printed or exported to another tool.

    When you first open the Access Rules page, the Generate Report button is replaced with a progress bar. After conflict analysis has completed, the Generate Report button becomes available along with the other conflict detection features.
  - Update the hit count information displayed in the table. For more information, see Viewing Hit Count Details, page 16-36 and Hit Count Selection Summary Dialog Box, page 16-20.
  - Run a policy query, which can help you evaluate your rules and identify ineffective rules. See Generating Policy Query Reports, page 12-28
  - Find and replace items within rules (button with the binoculars icon)—For more information, see Finding and Replacing Items in Rules Tables, page 12-16.
  - Move and rearrange rules (up and down arrows)—For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
  - Add rules to the table (+ icon)—For more information, see Adding and Removing Rules, page 12-9.
  - Edit the selected rule (pencil icon)—For more information, see Editing Rules, page 12-10.
  - Delete the selected rule (trash can icon)—For more information, see Adding and Removing Rules, page 12-9.
- **Conflict Navigation Bar (7)**—Use the Conflict navigation bar to navigate to conflicting rules in the rules table. For more information, see Using Automatic Conflict Detection, page 16-28.

Adding and Removing Rules

When you work with policies that use rules tables, like many of the firewall rules policies, you can add rules to the policy using several methods:

- **Add Row** button (+ icon)—Clicking the Add Row button beneath the table is the standard method to add a new rule. Clicking this button opens the dialog box for adding rules that is specific to that type of policy. If you select a row or section heading, the new rule is added after the selected row. Otherwise, it is added at the end of the appropriate scope (typically, the local scope).
- Right-click a row and select **Add Row**—This is equivalent to selecting a row and clicking the Add Row button.
- Copy and paste—If you want to create a new rule that is similar to an existing rule, you can select the rule, right-click and select **Copy**, then select the row after which you want to place the rule, right-click and select **Paste**. This creates a duplicate rule, which you can select and edit (see Editing Rules, page 12-10).
- Cut and paste—Cut and paste is similar to copy and paste, except you are deleting the existing rule when you select the **Cut** command. Instead of cut and paste, consider moving the rule (see Moving Rules and the Importance of Rule Order, page 12-19).

When you no longer need a rule, you can remove it by selecting the rule and clicking the **Delete Row** button (trash can icon).

Tip Rather than deleting a rule, consider first disabling the rule. By disabling a rule, you remove it from the device (when you redeploy the configuration) without removing it from Security Manager. Then, if you discover that you really needed that rule after all, you can simply enable it and redeploy the configuration. If you delete the rule, you would have to recreate it (there is no undo function). Thus, you might want to develop a policy of deleting rules only after they have been disabled for a certain amount of time. For more information, see Enabling and Disabling Rules, page 12-20.

Related Topics

- Using Rules Tables, page 12-8
- Using Sections to Organize Rules Tables, page 12-20

Editing Rules

To edit an existing rule in any of the rules policies that use rules tables, select the rule and click the **Edit Row** button, or right-click and select **Edit Row**. This allows you to edit all aspects of the selected rule.

Tip You cannot edit any aspect of an inherited rule from a local device rule policy. Edit inherited rules in Policy view.

For most rule tables, you can also edit specific attributes, or table cells, instead of editing the entire rule, using commands in the right-click menu.

The ability to edit a cell is limited by whether it makes sense to edit the content. For example, Inspection Rules have many limitations based on how the rule is configured:

- If you applied the rule to All Interfaces, you cannot edit source or destination addresses, the interface, or the direction of the rule.
- If you selected Default Inspection Traffic for the traffic match criteria (without selecting the option to limit inspection between source and destination), or Custom Destination Ports, you cannot edit source or destination addresses.
- If you selected Destination Address and Port (IOS), you cannot edit source addresses.

The following cell-level commands are available, although the ability to edit multiple rows is not supported in all policies that use rule tables:

- **Add <Attribute Type>**—When you select multiple rows and right-click a Source, User, Destination, Services, or Interface cell, you can select the Add command to append entries to the data currently in the selected cells. The Add command's full name includes the name of the attribute, for example, Add Source.
- **Edit <Attribute Type>**—Most attributes allow you to edit the content. Editing replaces the content of the cell. You can edit a single cell, or select multiple rows and edit the contents of the same type of cell in all rows at once. The Edit command's full name includes the name of the attribute, for example, Edit Interfaces.
- **Edit <Entry>**—In some cases, when you edit Source, User, Destination, Services, or Interfaces, you can select an entry in the cell and edit just that entry. For example, if the Sources cell contains three network/host objects and an IP address, you can select any of them and edit the entry. The edit command includes the name of the entry, for example, Edit HostObject.

- **Remove <Entry>**—In some cases, when you edit Source, User, Destination, Services, or Interfaces, you can select an entry in the cell and remove the entry. You cannot remove the last entry in the cell, because the rule would become invalid. The remove command includes the name of the entry, for example, Remove IP.
- **Create <Object Type> Object from Cell Contents**—In the Sources, User, Destinations, and Services cells, you can select the Create command to create a policy object of the appropriate type. You can also select an entry in the cell and create a policy object from just the selected item. The create command includes the policy object type you can create, and the name of the item that is the source for the object, either cell contents for everything in the cell, or the name of an entry if you selected one. When creating network/host objects, you are always creating network/host group objects.
- **Show <Attribute Type> Contents; Show <Entry> Contents**—The show commands let you view the actual data defined in the cell. The results depend on the view you are in:
  - Device View, Map View, or Import Rules—You are shown the actual IP addresses, fully-qualified domain names (FQDNs), services, or interfaces to which the rule will apply for the specific device. For example, if the rule uses network/host objects, you will see the specific IP addresses or FQDNs defined by the objects. If the rule uses interface objects, you will see the specific interfaces defined on the device that the object identifies, if any.

    The IP addresses for network/host objects are sorted in ascending order on the IP address, and then descending order on the subnet mask.

    Service objects are sorted on protocol, source port, and destination port.

    Interface objects are listed in alphabetical order. If the interface is selected because it matches a pattern in an interface object, the pattern is listed first, and the matching interface is shown in parentheses. For example, "* (Ethernet1)" indicates that the Ethernet1 interface on the device is selected because it matches the * pattern (which matches all interfaces).
- Policy View—You are shown the patterns defined in the policy objects and entries defined for the policy. Entries are sorted alphabetically, with numbers and special characters coming first.

Related Topics

- Using Rules Tables, page 12-8
- Adding and Removing Rules, page 12-9
- Moving Rules and the Importance of Rule Order, page 12-19
- Enabling and Disabling Rules, page 12-20
- Using Sections to Organize Rules Tables, page 12-20

Adding or Editing Address Cells in Rules Tables

Use the Add or Edit Sources or Destinations dialog boxes, or Address dialog boxes for NAT tables, to edit the source or destination entry in a rules table that includes sources or destinations. For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

You can enter any combination of the following address types to define the source or destination of the traffic. The type of policy determines whether an IPv4 or IPv6 address is required; you cannot mix address types. You can enter more than one value by separating the items with commas. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

- Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new objects from the selection list.

Note The only way to specify a fully-qualified domain name (FQDN) is to use an FQDN network/host object or a group object that includes an FQDN object. You cannot directly type in an FQDN. Not all policy types allow FQDN; you are prevented from specifying an object that contains an FQDN object if the policy does not allow it.

- Host IP address, for example, 10.10.10.100 (IPv4) or 2001:DB8::200C:417A (IPv6).
- IPv4 network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0.
- IPv6 network address and prefix length in the format 2001:DB8::/32.
- A range of IP addresses, for example, 10.10.10.100-10.10.10.200 (IPv4) or 2001:DB8::1-2001:DB8::100 (IPv6).
- (IPv4 only.) An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).
- Interface roles object. Enter the name of the object or click **Select** to select it from a list (you must select Interface Role as the object type). When you use an interface role, the rule behaves as if you supplied the IPv4 or IPv6 address of the selected interface. This is useful for interfaces that get their address through DHCP, because you do not know what IP address will be assigned to the device. For more information, see Understanding Interface Role Objects, page 6-73.

  If you select an interface role as a source, the dialog box displays tabs to differentiate between hosts or networks and interface roles.

Navigation Path

Do any of the following in a rules policy that includes sources, destinations, or other address cells:

- Right-click an address cell in a rules table and select **Edit Sources** or **Edit Destinations** or a similar command. The data replaces the content of the selected cells.
- Select an entry in an address cell and select **Edit <Entry>**. The data replaces the selected entry.
- Select multiple rules, right-click a Sources or Destination cell, and select **Add Sources** or **Add Destinations**. The data is appended to the data already in the cell.

Adding or Editing User Cells in Rules Tables

Tip The user cell applies to ASA 8.4(2+) only. Anything configured in the cell is ignored for other device types or OS versions.

Use the Add or Edit Users dialog boxes to edit the user entry in a rules table that includes user identity groups. For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

You can enter any combination of the following to identify traffic based on Active Directory (AD) user or user group names. If you configure identity user groups, they apply to source traffic only. For traffic to match the rule, both the source addresses and identity user groups must match. That is, the rule applies to traffic sent from users on the specific networks or hosts defined in the source field when directed at the destination. For more information, see Configuring Identity-Based Firewall Rules, page 13-21.

To make the rule apply to a user without regard for the source address, specify **any** in the source cell.

You can enter more than one value by separating the items with commas. Following are the supported formats:

- Identity user group objects.
- Individual users: NETBIOS_DOMAIN\user
- User groups (note the double \): NETBIOS_DOMAIN\\user_group

Click **Select** to select objects, users, or user groups from a list or to create new objects. For more information, see Selecting Identity Users in Policies, page 13-21 and Creating Identity User Group Objects, page 13-19.

Navigation Path

Do any of the following in a rules policy that includes user cells:

- Right-click a user cell in a rules table and select **Edit Users**. The data replaces the content of the selected cells.
- Select an entry in a user cell and select **Edit <Entry>**. The data replaces the selected entry.
- Select multiple rules, right-click a user cell, and select **Add Users**. The data is appended to the data already in the cell.

Adding or Editing Services Cells in Rules Tables

Use the Edit Services dialog box to edit the services that define the type of traffic to act on. You can enter more than one value by separating the items with commas.

You can enter any combination of service objects and service types (which are typically a protocol and port combination). If you type in a service, you are prompted as you type with valid values. You can select a value from the list and press Enter or Tab. You can also click **Select** to select the service from a list, or to create a new service.

For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.

For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

Navigation Path

Do any of the following in a rules policy that includes services:

- Right-click a Services cell in a rules table and select **Edit Services**. The data replaces the content of the selected cells.
- Select an entry in a Services cell and select **Edit <Entry>**. The data replaces the selected entry.
- Select multiple rules, right-click a Services cell, and select **Add Services**. The data is appended to the data already in the cell.

Tip For inspection rules, services appear in the Traffic Match column and only for rules where the traffic matches source, destination, and port.

Adding or Editing Interfaces or Zones Cells in Rules Tables

Use the Add or Edit Interfaces (or Zones) dialog box to edit the interfaces or zones for which the rule is defined. For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

- When editing interfaces, you can enter any combination of specific interface names or interface roles. You can enter more than one value by separating the items with commas. Enter the names or click **Select** to select the interfaces and roles from a list, or to create new roles. An interface must already be defined to appear on the list.

  When you deploy the policy to the device, interface roles are replaced by actual interface names, and only to interfaces that are actually configured on the device. To see which interfaces will actually be selected by a rule, right-click the Interfaces cell and select **Show Interfaces**.
- When editing zones, you can select only one interface role, and you cannot select individual interfaces. The interface roles are used to create zones for zone based firewall rules. To see the interfaces that will belong to the zone, right-click the Zones cell and select **Show Zone Contents**.

For more information about interface roles and selecting interfaces, see the following topics:

- Understanding Interface Role Objects, page 6-73
- Specifying Interfaces During Policy Definition, page 6-76

Navigation Path

Do any of the following in a rules policy that includes interfaces or zones:

- Right-click an Interfaces or Zones cell in a rules table and select **Edit Interfaces**, **Edit Zones**, or similar command. The data replaces the content of the selected cells.
- Select an entry in an Interfaces cell and select **Edit <Entry>**. The data replaces the selected entry. You cannot edit an entry in a zone.
- Select multiple rules, right-click an Interfaces cell, and select **Add Interfaces**. The data is appended to the data already in the cell. You cannot add entries to a zone.

Editing Category Cells in Rules Tables

Use the Edit Category dialog box to change the category assigned to a rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13. For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

Navigation Path

Right-click a Category cell in a rules policy that includes categories and select **Edit Category**.

Editing Description Cells in Rules Tables

Use the Edit Description dialog box to edit the description of the rule. The description helps you identify the purpose of a rule and can be up to 1024 characters. For detailed information on editing rules cells, see Editing Rules, page 12-10.

Navigation Path

Right-click a Description cell in a rules policy that includes descriptions and select **Edit Description**.

Showing the Contents of Cells in Rules Tables

Use the Show Contents dialog boxes to display the actual, translated data defined in a source, user, destination, services, interfaces, zones, or other cell in a rules table that includes addresses, identity user groups, interfaces, services, or policy objects that define those things. The title of the dialog box

indicates which cell or entry you are examining. Use this information to determine to which addresses, services, or interfaces the rule will actually apply when deployed to the device. For detailed information about editing or viewing cell contents, see Editing Rules, page 12-10.

What you see in the dialog box depends on the view you are in:

- Device View, Map View—You are shown the actual IP addresses, users, services, or interfaces to which the rule will apply for the specific device. For example, if the rule uses network/host objects, you will see the specific IP addresses or fully-qualified domain names (FQDN) defined by the objects. If the rule uses interface objects, you will see the specific interfaces defined on the device that the object identifies, if any.
  - The IP addresses for network/host objects are sorted in ascending order on the IP address, and then descending order on the subnet mask.
  - Service objects are sorted on protocol, source port, and destination port.
  - Interface objects are listed in alphabetical order. If the interface is selected because it matches a pattern in an interface object, the pattern is listed first, and the matching interface is shown in parentheses. For example, "* (Ethernet1)" indicates that the Ethernet1 interface on the device is selected because it matches the * pattern (which matches all interfaces).
- Policy View—You are shown the patterns defined in the policy objects and entries defined for the policy. Entries are sorted alphabetically, with numbers and special characters coming first.

Filtering Contents

A List Filter field is provided above the results in the Show Contents dialog box. You can use the List Filter field to quickly locate any entries that contain a specified text string.

Figure 12-2 List Filter Field

1	Filter-parameters button.	2	Clear button.

To search for a specific text string in the Show Contents list:

- Click in the List Filter field to place the text cursor, and then begin typing.

  These are "live filter" fields. That is, as you type each character, entries that do not include your current text string are removed from the list or table.

To clear a List Filter field:

- Click the clear button at the right side of the field.

  This button appears when you begin typing in the field. (You also can highlight the characters and press the Delete or Backspace key on your keyboard.)

  When you clear the List Filter field, all entries in the list are again displayed.

You can tune the filter results by selecting case sensitivity or insensitivity, by allowing wildcards or regular expressions, and by specifying where in a returned string your characters must be located.

To change the List Filter criteria:

1. Click the filter-parameters button (magnifying glass) at the left side of the List Filter field to open the parameters menu.
2. Choose an option.

   The menu consists of three sections:

   – **Case sensitive** and **Case insensitive** – Choose one or the other. If you choose **Case sensitive**, found text must match not only the characters you enter, but also their as-typed case.

   – **Use wildcards** and **Use regular expression** – Choose one or the other. The following wildcards are recognized:

   – ***** (asterisk) – Match zero or more characters at that location in the string.

   – **+** (plus sign)– Match one or more characters at that location in the string.

   – **?** (question mark) – Match one character at that location in the string.

   – **Match from start**, **Match exactly**, and **Match anywhere** – Choose one. **Match from start** means that the string you enter must be found at the beginning of an entry, although it can be part of a larger set of characters. **Match exactly** requires that the string you enter exactly match the entire column entry. **Match anywhere** means the string can be found anywhere within an entry, and it can be part of a larger set of characters.
3. Repeat Steps 1 and 2 to change another parameter.

Navigation Path

Do any of the following in a rules policy that includes sources, user, destinations, services, interfaces, zones, or other fields that specify networks, identity user groups, interfaces, or services. You can also show contents when using tools that work with rules, such as importing rules.

- Right-click one of those cells and select **Show <Attribute Type> Contents**, where the attribute type is the name of the cell. The data includes all entries defined in the cell.
- Right-click an entry in one of those cells and select **Show <Entry> Contents**, where the name of the selected entry is included in the command name. The data displayed is only for the selected entry.

Tip For inspection rules, services appear in the Traffic Match column and only for rules where the traffic matches source, destination, and port.

Finding and Replacing Items in Rules Tables

In policies that use rules tables, you can search for items in some cells and selectively replace them. The cells that you can search depend on the policy. You can use wildcard characters to find items based on pattern matching, for example, so that you can replace several related networks with a new network/host policy object defined for them.

To use find and replace, click the **Find and Replace** (binoculars icon) button at the bottom of any policy that uses rules tables to open the Find and Replace Dialog Box, page 12-17. In the Firewall folder, this includes AAA rules, access rules, IPv6 access rules, inspection rules, zone based firewall rules, and web filter rules (for ASA/PIX/FWSM devices only). For ASA/PIX/FWSM devices, it also includes the NAT translation rules policy (but not for every combination of context and operational mode) and the IOS, QoS, and connection rules platform service policy.

When searching for items, you select the type of item, the columns you want to search, and enter the string that you want to find and optionally, the string you want to use to replace it. You can find and replace the following types of items:

- Network—A network/host object name, or the IP address of a host or network.
- User—An Active Directory (AD) username (NetBIOS_DOMAIN\user), user group name (NetBIOS_DOMAIN\\user_group), or identity user group object name.
- Service—A service object name or protocol and port, for example TCP/80. The search is syntactic, not semantic, that is, if you are searching for TCP/80 and a rule uses HTTP, the search results will not find it.
- Interface Role—An interface name or interface role object name.

Note In access rules, you can search for global rules by using the Global interface name. However, there is no way to convert between global and interface-specific rules. Although you can find global rules using the Global interface name, if you try to replace an interface name with the name "Global," you are actually creating an interface-specific access rule that uses a policy object named Global.

- Text—A text string in a Description field.

The following are some examples of what you might do with find and replace:

- If you create a new network/host object named network10.100 for all networks in the 10.100.0.0/16 range, you can search and replace all subordinate network specifications. For example, you can search for ^10.100* to find all addresses like 10.100.10.0/24. Select the **Find Whole Words Only** and **Allow Wildcard** options, and enter network10.100 as the replacement string. Because you selected Find Whole Words Only, the string that is replaced is the entire 10.100.10.0/24 string, not just the 10.100 portion.
- If you want to find all rules that use IP addresses (instead of network/host objects), you can search for *.*.*.* to find all host or network IP addresses. You can then selectively edit the cell while the Find and Replace dialog box is open.
- If you want to replace all interface role objects that include "side" in the name (such as inside and outside) with the interface role object named External, search for *side with the **Find Whole Words Only** and **Allow Wildcard** options selected, and enter External in the Replace field.

Related Topics

- Editing Rules, page 12-10

Find and Replace Dialog Box

Use the Find and Replace dialog box to locate and optionally replace items in rule table cells. The types of items you can search for differ based on the policy you are viewing.

Navigation Path

Click the **Find and Replace** (binoculars icon) button at the bottom of any policy that uses rules tables. In the Firewall folder, this includes AAA rules, access rules, IPv6 access rules, inspection rules, zone based firewall rules, and web filter rules (for ASA/PIX/FWSM devices only). For ASA/PIX/FWSM devices, it also includes the NAT translation rules policy (but not for every combination of context and operational mode) and the IOS, QoS, and connection rules platform service policy.

Related Topics

- Finding and Replacing Items in Rules Tables, page 12-16
- Editing Rules, page 12-10

Field Reference

Table 12-2 Find and Replace Page

Element	Description
Type	The type of item you are trying to find. Select the type, then select which columns you want to search. If you select All Columns, the columns searched are those also listed with the All Columns item (the search does not consider every column in the table). • Network—A network/host object name, or the IP address of a host or network. • User—An Active Directory (AD) username (NetBIOS_DOMAIN\user), user group name (NetBIOS_DOMAIN\\user_group), or identity user group object name. • Service—A service object name or protocol and port, for example TCP/80. The search is syntactic, not semantic, that is, if you are searching for TCP/80 and a rule uses HTTP, the search results will not find it. • Interface Role—An interface name or interface role object name. **Note** In access rules, you can search for global rules by using the Global interface name. However, there is no way to convert between global and interface-specific rules. Although you can find global rules using the Global interface name, if you try to replace an interface name with the name "Global," you are actually creating an interface-specific access rule that uses a policy object named Global. • Text—A text string in a Description field.
Find	The string you are trying to locate. If you are searching for a policy object, click **Select** to choose the object from a list.
Replace	(Optional) The string you want to use to replace the search string. What gets replaced is controlled by the search options. If you want to replace the search string with the name of a policy object, click **Select** to choose the object from a list. You can replace search strings with multiple items. Separate the items with commas. For example, you can search for the TCP service and replace it with TCP, UDP. You can remove items by not entering anything in the Replace field and clicking the **Replace** button. This field is greyed out if the table does not allow editing.
Direction	The direction in which you want to search relative to the currently selected row or cell, either up or down. When the end of the table is reach, the search continues to the top of the table.

Table 12-2 Find and Replace Page (continued)

Element	Description
Match Case	For text searches, whether you want to match the capitalization you used in the Find field.
Find Whole Words Only	Whether the search should find and select only whole words, which are strings delimited by spaces or punctuation. For example, a whole word search for SanJose will find SanJose but not SanJose1. If you use this option with the Allow Wildcard option, you can search for partial strings but if you replace the located string, you replace the whole word and not the partial string. For example, you can search for ^10.100* to find all addresses like 10.100.10.0/24, and replace with them with the network10.100 policy object. By selecting Whole Words, the network/host object replaces the entire address, not just the portion you searched for. For text searches, this option and the Allow Wildcards option are mutually exclusive.
Allow Wildcards	Whether the search or replacement strings use wildcard characters. If you do not select this option, all characters are treated literally. You can use the Java regular expression syntax to create your expression with the following exceptions: • Period (.)—The period is a literal period and it is implicitly escaped. • Question mark (?)—The question mark indicates a single character. • Asterisk (*)—The asterisk matches one or more characters. It does not match zero characters. • Plus sign (+)—The plus sign means the same as the asterisk; it matches one or more characters.
Find Next button	Click this button to find the next occurrence of the search string.
Replace button	Click this button to replace the found string with the replacement string.
Replace All button	Click this button to automatically find the search string and replace it throughout the table.

Moving Rules and the Importance of Rule Order

Rules policies that use rules tables are ordered lists. That is, the top to bottom order of the rules matters and has an effect on the policy.

When the device analyzes a packet against a rules policy, the device searches the rules in order from top to bottom. The first rule that matches the packet is the rule that is applied to the packet, and all subsequent rules are ignored. Thus, if you place a general rule pertaining to IP traffic before a more specific rule pertaining to HTML traffic for a given source or destination, the more specific rule might never be applied.

For access control rules, you can use the automatic conflict detection tool to help identify when rule order will prevent a rule from ever being applied to traffic (for more information, see Using Automatic Conflict Detection, page 16-28). For other rules policies, carefully inspect the table to spot problems with rule order.

When you find that you need to rearrange the order of a rule, select the rule that needs to be moved and click the **Up Row** (up arrow) or **Down Row** (down arrow) buttons as appropriate. If these buttons do not appear beneath the rules table, rule order does not matter and you cannot rearrange them.

If you use sections to organize your rules, you can move rules only within the section. When you move rules that are outside the sections, you can move them above or below the section. For more information about working with sections, see Using Sections to Organize Rules Tables, page 12-20.

Tip Special rules apply to moving access rules when you mix interface-specific and global rules in a policy. For more information, see Understanding Global Access Rules, page 16-3.

Related Topics

- Using Rules Tables, page 12-8
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Using Sections to Organize Rules Tables, page 12-20

Enabling and Disabling Rules

You can enable and disable individual rules in a policy that uses rules tables, such as most firewall services rules policies. Your change takes effect when you redeploy the configuration to the device.

If a rule is disabled, it appears in the table overlain with hash marks. When you deploy the configuration, disabled rules are removed from the device.

Disabled rules are kept in the rules policy in Security Manager as a convenience so that you can easily enable needed rules without recreating them. Thus, it is often wise to disable a rule that you believe you no longer need instead of immediately deleting it.

To change whether a rule is enabled or disabled, select the rule, right-click and select **Enable** or **Disable**, as appropriate.

Related Topics

- Using Rules Tables, page 12-8
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Moving Rules and the Importance of Rule Order, page 12-19
- Using Sections to Organize Rules Tables, page 12-20

Using Sections to Organize Rules Tables

You can organize policies that use rules tables into sections. There are two types of sections:

- Scopes, which define inheritance relationships between a policy and an inherited policy. These sections are automatically created when you inherit policies. For more information, see Understanding Rule Inheritance, page 5-4.
- User-defined sections, which are convenient groupings that help you organize rules so that you can evaluate and edit the policy more easily. These types of sections are most useful for policies that contain a large number of rules.

  All rules within a section must be sequential; you cannot group rules randomly. If you want to identify non-contiguous rules as being related, you can assign the same category to the rules.

User-defined sections are set off visually from the other rules in the table by an indented section heading. The heading contains, left to right, a +/- icon for opening and closing the section, a band of color identifying the category you assigned the section (if any), the section name, the first and last rule number contained in the section (for example, 4-8), and the description, if any, you gave the section.

Note You might need to resize the rule number column to see the numbering of rules in sections.

Whether you create user-defined sections is completely up to you. If you decide creating these types of sections is worthwhile, the following information explains how to create and use them:

- To create a new section, right-click a row that you want to place the section and select **Include in New Section**. (You can also use Shift+click to select a block of rules.) You are prompted for a name, description, and category for the section (the name is the only required element).
- To move existing rules into a section, select one or more contiguous rules, right-click and select **Include in Section <name of section>**. This command appears only if the selected rows are next to an existing section. If the rows you want to add to a section are not currently next to the section, you can do one of two things: move the rules until they are next to the section; cut the rules and paste them into the section.
- You cannot move a section. Instead, you need to move the rules that are outside of the section around it. When you move a rule that is next to, but not within, a section, the rule jumps over the section.
- You cannot move a rule outside of or through a section. A section defines the borders within which you can move rules. If you want to move rules out of a section and back into the Local scope section, select one or more contiguous rules, right-click and select **Remove from Section <name of section>**. The rules must be at the beginning or end of the section to use this command. If they are not, you can either move the rules until they are, or use cut and paste to move them out of the section.
- To add a new rule to a section, select the rule after which you want to create the rule before clicking the Add Row button. To place it at the beginning of the section, select the section heading.

  If you want to create a rule after, but outside of, a section, you can either create it as the last rule in the section and then remove it from the section, or create it just above the section and click the down arrow button.
- You can change the name, description, or category of a section by right-clicking the section heading and selecting **Edit Section**.
- When you delete a section, all rules contained in the section are retained and moved back into the Local scope section. No rules are deleted. To delete a section, right-click the section heading and select **Delete Section**.
- If you use the Combine Rules tool, the resulting combined rules respect your sections. Rules that are in a section can be combined only with other rules in that section.

Related Topics

- Using Rules Tables, page 12-8

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Moving Rules and the Importance of Rule Order, page 12-19
- Enabling and Disabling Rules, page 12-20

Add and Edit Rule Section Dialog Boxes

Use the Add and Edit Rule Section dialog boxes to add or edit a user-defined section heading in a rules table. For detailed information about how to use sections to organize a rules table, see Using Sections to Organize Rules Tables, page 12-20.

Navigation Path

Do one of the following:

- Select one or more rules in a rules table, right-click and select **Include in New Section**.
- Right-click a section heading and select **Edit Section**.

Field Reference

Table 12-3 Add and Edit Rule Section Dialog Boxes

Element	Description
Name	The name of the section.
Description	A description for the section, up to 1024 characters.
Category	The category assigned to the section. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Combining Rules

Access rules and AAA rules policies can grow over time to include a large number of rules. The size of these policies can make it difficult to manage them. To alleviate this problem, you can use the rule combiner tool to reduce the number of rules in a policy without changing how the policy handles traffic.

Tip Combining rules can dramatically compress the number of access rules required to implement a particular security policy. For example, a policy that required 3,300 access rules might only require 40 rules after hosts and services are efficiently grouped. However, you cannot use the rule combiner with IPv6 access rules or with rules that specify users or user groups, either directly or with identity user group objects. You can use the tool with rules that use FQDN network/host objects.

You might have several rules that allow a specific range of services to various trusted hosts (as sources) to various public servers (as destinations). If you have 10 rules applying to this situation, it is possible that those 10 rules can be combined into a single rule. You could then create new policy objects for the collection of services (for example, AllowedServices), hosts (for example, TrustedHosts), and servers (for example, PublicServers). To create the new objects during rule combination, you can right-click the newly-combined cells and select **Create Network (or Service) Object from Cell Contents**.

For example, you might have two rules for interface FastEthernet0:

- Permit TCP for source 10.100.10.1 to destination 10.100.12.1

- Permit TCP for source 10.100.10.1 to destination 10.100.13.1

These can be combined into a single rule: permit TCP for source 10.100.10.1 to destination 10.100.12.1, 10.100.13.1.

Multidimensional sorting is used to combine rules. For example, for access rules:

1. Rules are sorted by their sources, so rules with the same source are placed together.
2. Same-source rules are sorted by destination, so rules with the same source and destination are placed together.
3. Same-source and same-destination rules are combined into a single rule, and the services are concatenated.
4. Adjacent rules are checked to see if they have the same source and service. If so, they are combined into a single rule, and the destinations are concatenated.
5. Adjacent rules are checked to see if they have the same destination and service. If so, they are combined into a single rule, and the sources are concatenated.

Sorting is repeated based on destination and service in place of source.

Tip Rules from different sections are never combined. Any sections you create to organize rules limit the scope of the possible combinations. Also, interface-specific and global access rules are never combined. For more information about global rules, see Understanding Global Access Rules, page 16-3.

Related Topics

- Chapter 15, "Managing Firewall AAA Rules"
- Chapter 16, "Managing Firewall Access Rules"

Step 1 Select the policy whose rules you want to combine from the **Firewall** folder. You can combine rules for the following types of policy:

- AAA rules
- Access rules

Step 2 If you want the tool to limit possible combinations to a specific group of rules, select them. You can select rules using Shift+click and Ctrl+click, select all rules in a section by selecting the section heading, or all rules within a scope by selecting the scope heading (for example, Local). To not limit the tool, do not select anything in the table. Keep the following in mind:

- In Device view, you can save combinations only for local rules. The tool will allow you to run it on shared and inherited rules, but you cannot save the results. If you do not select any rules, the default is to consider all local scope rules.
- To combine rules in shared policies, you must run the tool in Policy view. If you do not select any rules, the default is to consider all mandatory rules.

You are warned if you try to run the tool when you cannot save the results.

Step 3 Right-click anywhere in the rules table and choose **Combine Rules** to open the Combine Rules Selection Summary Dialog Box, page 12-24. If you have selected specific rules for which to limit combining, make sure you right-click on one of the selected rules or the rules will be deselected.

Step 4 Select the columns you want the rule to consider combining. If you do not select certain columns, the combined rules must have the identical settings in those columns to be combined.

You can also elect to consider combining the rules you selected or all rules within the policy.

Tip If a column type is not listed, then combined rules must have the identical content in those cells except for the Description cell. Rules that have different content for the cells are not combined.

Step 5 Click **OK** to generate the combination and display the results in the Rule Combiner Results Dialog Box.

Analyze the results and evaluate whether you want to save the combinations. You must save all or none, you cannot pick and choose which combinations to save.

For more information on evaluating the results, see Interpreting Rule Combiner Results, page 12-25. For an example, see Example Rule Combiner Results, page 12-27.

Step 6 Click **OK** to replace the original rules in the rules tables with the combined rules.

Combine Rules Selection Summary Dialog Box

Use the Combine Rules Selection Summary dialog box to define the parameters used for combining rules in firewall rules policies. When you click **OK**, the combination results are displayed in the Rule Combiner Results Dialog Box, where you can choose to save or discard the results as explained in Interpreting Rule Combiner Results, page 12-25.

Navigation Path

You can combine rules from the AAA Rules Page, page 15-10 and the Access Rules Page, page 16-10. Click **Tools** located at the bottom of the tables and select **Combine Rules**.

Field Reference

Table 12-4 Combine Rules Selection Summary Dialog Box

Element	Description
Policy Selected	Shows the policy selected and the scope. Local indicates the local device rules. Otherwise, the field indicates the name of the shared policy and the scope selected within the policy, if any.
Rules to be combined	The rules you want the tool to consider combining: • All Rules—Consider combining all rules within the selected policy. • Selected Rules—Consider combining only those rules you selected in the policy before starting the tool. For detailed information on selecting rules before running the tool, see Combining Rules, page 12-22.

Table 12-4 Combine Rules Selection Summary Dialog Box (continued)

Element	Description
Choose which columns to combine	The columns in the rules table that can be combined. Any columns that you do not select must have the identical content for two rules to be combined (even those not listed as combinable, except for the Description column). The columns you can combine are: • Source • User • Destination • Service • Interface • Security Sources • Security Destinations • For AAA rules, these additional columns: – Action – Auth Proxy

Interpreting Rule Combiner Results

Use the Rule Combiner Results dialog box to evaluate the results of a rule combination (see Combining Rules, page 12-22). The dialog box includes a summary of the results, and shows the new rules that will be created if you click **OK**.

Changed rule cells are outlined in red. Select a combined rule in the upper table to see the rules in the lower table that were combined to create the rule.

You can refine some elements of the results in this window:

- You can right-click on the Source, Destination, and Service cells with multiple elements and select **Create Network (or Service) Object from Cell Contents** to create a new policy object that contains the contents of the combined cell. The new object replaces the contents of the cell.

  You can also automatically create network object groups in the deployed configuration to replace the comma-separated values in a rule table cell. The network objects are created during deployment, and they do not affect the content of your rules policy. To enable this option, select **Tools > Security Manager Administration > Deployment** to open the Deployment Page, page 11-13 and select **Create Object Groups for Multiple Sources, Destinations, or Services in a Rule**.
- You can right-click on Description and select **Edit Description** to change the description. The descriptions of combined rules are a concatenation of the descriptions of the old rules separated by new lines.

For an example, see Example Rule Combiner Results, page 12-27.

Tips

- The combined results are not applied to the policy until you click **OK**. If you do not like the results of the combination, click **Cancel** and consider selecting smaller groups of rules to limit the scope of the Combine Rules tool.

If you click **OK** but then decide you do not want to accept the changes, you have two options. First, make sure you do not click **Save** on the policy page, select a different policy, and click **No** when prompted to save your changes to the policy. If you already clicked **Save**, you can still back out the changes by discarding your activity or configuration session (for example, **File > Discard** in non-Workflow mode), but this also discards any other changes you have made to other policies. Once you submit your changes or your activity is approved, you cannot undo your changes.

- You are allowed to run the Combine Rules tool even if you are combining rules for a policy that you are not allowed to save. For example, you cannot save combined rules for a shared or inherited policy in Device view. You are warned before running the tool if you will not be allowed to save the results.
- Rules from different sections are never combined. Any sections you create to organize rules limit the scope of the possible combinations. Also, interface-specific and global access rules are never combined. For more information about global rules, see Understanding Global Access Rules, page 16-3.

Navigation Path

You can combine rules from the AAA Rules Page and the Access Rules Page. Click **Tools** located at the bottom of the tables and select **Combine Rules**, fill in the Combine Rules Selection Summary Dialog Box and click **OK**.

Field Reference

Table 12-5 Combined Rules Results Summary

Element	Description
Result Summary	Provides a summary of the results of the combination and indicates the number of original rules, the number of rules remaining after the combination, and the number of changed and unchanged rules, if any combinations could be made.
Resulting Rules table	The rules that will replace the rules currently in the policy. If you click **OK**, these rules become part of your policy. The columns are the same as those in the associated policy (see AAA Rules Page, page 15-10 or Access Rules Page, page 16-10), with the addition of the Rule State column. The Rule State column shows the status of the rule: • Modified, Combined—The new rule is the result of combining one or more rules or modifying an existing rule. A red box around a cell indicates cells that have combined contents. • Unchanged—The rule remains unchanged, as it could not be combined with any other rule. • Not Selected—You did not select the rule for possible combination. If there are a large number of rules, you can use the buttons beneath the table to scroll through the rules that have changes. Unchanged and unselected rules are skipped.
Original rules table (lower table)	The table in the lower half of the dialog box shows the original rules that were combined to create the rule you select in the upper table.

Table 12-5 Combined Rules Results Summary (continued)

Element	Description
Detail Report button	Click this button to create an HTML report of the results. The report summarizes the results and also provides the details about the resulting rules and the rules that were combined to create the new rule. For combined rules that have a lot of entries in cells, this report makes it easier to read the results. You can also print or save the report for later use.

Example Rule Combiner Results

When you run the Combine Rules tool as described in Combining Rules, page 12-22, the results of the combination are displayed in the Rule Combiner Results Dialog Box (see Interpreting Rule Combiner Results, page 12-25).

Figure 12-3 shows an example of a rule combination.

The new rules are shown in the upper table. Any new rules are indicated as modified or combined rules, and the changed cells are outlined in red. When you select a new rule in the upper table, the lower table shows the old rules that were combined to create the new rule. In this example, the two old rules had the same destination, service, and interface, and the two distinct sources were concatenated to form the new rule.

The top of the report summarizes the results. In this example, 5700 rules were reduced to 96 rules.

Figure 12-3 Example of Rule Combiner Results

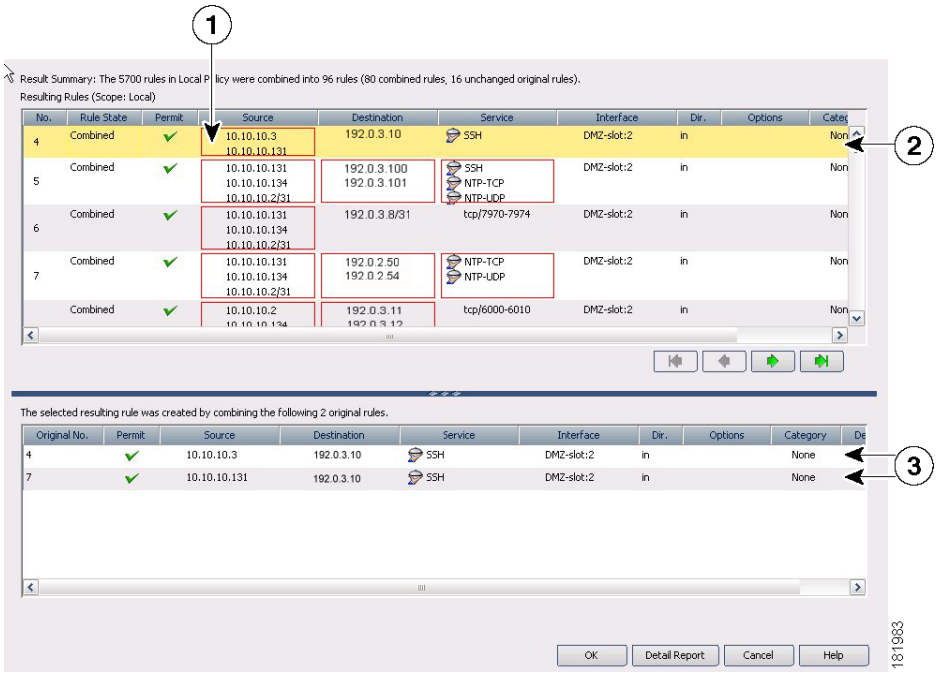

1	Combined cell	**3**	Original rules
2	Newly combined rule		

Related Topics

- Chapter 15, "Managing Firewall AAA Rules"
- Chapter 16, "Managing Firewall Access Rules"

Converting IPv4 Rules to Unified Rules

Prior to the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, separate pages, policies and policy objects were provided for configuring IPv4 and IPv6 firewall rules and policies. With Security Manager 4.4 and ASA 9.0+, these policies and policy objects were combined or unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided in Device view, while in Policy view, IPv4 and unified versions of the AAA-, access- and inspection-rule policy types are provided.

A utility to convert separate IPv4 and IPv6 firewall rules to "unified" rules is provided with Security Manager 4.4 for use when you upgrade an ASA from an earlier version to 9.0 or later.

Navigation Path

To access the firewall-rule unification utility:

- (Policy view) Select the firewall IPv4 rule type from the Policy Type selector and then right-click the desired policy in the Policies pane; choose **Convert to <*rule-type*>** Rules (Unified).

Related Topics

- AAA Rules Page, page 15-10
- Access Rules Page, page 16-10
- Inspection Rules Page, page 17-8

Step 1 Open the utility as described above; in the Convert Policy dialog box, provide a name for the new unified policy and click OK.

Following processing, the new unified rules policy is displayed. You can now assign this policy to ASA 9.0+ devices.

Generating Policy Query Reports

For most of the firewall rules policies, you can generate policy query reports that can help you evaluate your rules. With policy query reports, you can determine what rules already exist for a particular source, user, destination, interface, service, or zone before creating new rules to apply to those items.

To a limited degree, you can also determine if there are some blocking rules that prevent a rule from being used, or redundant rules that you can delete. If you are evaluating access rules, however, it is better to use the more powerful rule analysis tool to determine these problems.

When you create a policy query, you describe the traffic that interests you, much the same way you describe traffic when creating a rule. Creating a query is essentially the same as creating a rule, but you might want to describe the rule more broadly to capture a wider set of traffic so you can see a set of related rules rather than a single rule or a limited number of rules. The query you create depends on the information you are trying to discover.

The possible extent of a query depends on the view you are in:

- Device or Map view—The query is limited to the selected device. However, you can query across all supported rule types. This allows you to compare different types of rules that apply to the same traffic.
- Policy view—The query is limited to the selected policy. You see only rules that are defined in that policy, and you cannot query other types of policies. If you want to query a shared policy while examining other policies, select a device that is assigned to the shared policy, and query the policy from the device in Device view.

Related Topics

- AAA Rules Page, page 15-10
- Access Rules Page, page 16-10
- Inspection Rules Page, page 17-8
- Web Filter Rules Page (ASA/PIX/FWSM), page 18-3
- Zone-based Firewall Rules Page, page 21-58

Step 1 Select the policy that you want to query from the **Firewall** folder. You can query any of the following types of policy:

- AAA Rules
- Access Rules
- Inspection Rules
- Web Filter Rules (PIX/ASA/FWSM)
- Zone Based Rules

Step 2 Click the **Query** button located below the table to open the Querying Device or Policy dialog box.

Step 3 Enter the parameters that define the rules you want to query. When setting up your query, you must select at least one rule type; enabled, disabled or both; permitted, denied, or both; and mandatory, default, or both. For detailed information about the query parameters, see Querying Device or Policy Dialog Box, page 12-29.

In Policy view, you cannot change the type of rule you are querying. In Device view, you can query any combination of rule types.

Step 4 Click **OK** to view the rules that match the criteria in the Policy Query Results dialog box. For information on reading the report, see Interpreting Policy Query Results, page 12-32.

For an example of a policy query report, see Example Policy Query Result, page 12-34.

Querying Device or Policy Dialog Box

Use the Querying Device or Querying Policy dialog box to set up the parameters for a query. The query results show the rules that match your parameters. The title of the dialog box indicates what you are querying:

- In Device or Map view, you are querying rules defined for the selected device.
- In Policy view, you are querying rules within the selected policy only.

You can query rules from these types of policies: AAA rules, access rules, inspection rules, web filter rules for ASA/PIX/FWSM, and zone based firewall rules.

When setting up your query, you must select at least one rule type; enabled, disabled or both; permitted, denied, or both; and mandatory, default, or both.

Note For inspection rules, if you enter Global as the interface value, the match status results will be shown as a partial match even if the match is complete.

Results are displayed in the Policy Query Results dialog box (see Interpreting Policy Query Results, page 12-32).

Navigation Path

To generate Policy Query reports, do one of the following:

- (Device view) Select a device, then select one of the supported firewall rules policies from the Firewall folder, and then click the **Query** button located below the table.
- (Policy view) Select any of the supported firewall rules policies from the Firewall folder, then select a specific policy from the Shared Policy selector, and then click the **Query** button located below the table.
- (Map view) Right-click a device and select a supported firewall rules policy from the Edit Firewall Policies menu. Click the **Query** button.

Related Topics

- Generating Policy Query Reports, page 12-28
- Example Policy Query Result, page 12-34

Field Reference

Table 12-6 Querying Device or Policy Dialog Box

Element	Description
Rule Types	The type of rules you want to query. When querying in Policy view, you cannot change the selection. When querying in Device view, you can select any of the following types of rules; the scope of the query is limited to the selected device: • AAA Rules • Access Rules • Inspection Rules • Web Filter Rules • Zone Based Rules
Enabled and/or Disabled Rules	Whether you want to query enabled or disabled rules, or both.
Mandatory and/or Default Rules	Whether you want to query rules that are in the mandatory or default sections, or both.
Match	Whether you want to query rules that permit or deny traffic, or both.

Table 12-6 Querying Device or Policy Dialog Box (continued)

Element	Description
Sources Destinations	The source or destination of the traffic. You can enter more than one value by separating the items with commas. **Note** If you leave a field blank, the query matches any address for that field. You can enter any combination of the following address types to define the source or destination of the traffic. For more information, see Specifying IP Addresses During Policy Definition, page 6-87. • Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list. • Host IP address, for example, 10.10.10.100. • Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0. • A range of IP addresses, for example, 10.10.10.100-10.10.10.200. • An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81). **Tip** You can create an object with a list of the IP addresses to facilitate future policy query requests.
User	(ASA 8.4(2+) only.) The Active Directory (AD) usernames, user groups, or identity user group objects for the rule, if any. You can enter more than one value by separating the items with commas. **Note** If you leave a field blank, the query matches only those rules that have nothing in the User field. You can enter any combination of the following values. • Individual user names: NetBIOS_DOMAIN\username • User groups (note the double \): NetBIOS_DOMAIN\\user_group • Identity user group object names. Click **Select** to select objects, users, or user groups from a list or to create new objects. For more information, see: • Selecting Identity Users in Policies, page 13-21 • Configuring Identity-Based Firewall Rules, page 13-21 • Creating Identity User Group Objects, page 13-19

Table 12-6 Querying Device or Policy Dialog Box (continued)

Element	Description
Services	The services that define the type of traffic that is acted on. You can enter more than one value by separating the items with commas. **Note** If you leave the field blank, the query matches any service. You can enter any combination of service objects and service types (which are typically a protocol and port combination). If you type in a service, you are prompted as you type with valid values. You can select a value from the list and press Enter or Tab. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100. **Tip** You can create an object with a list of the services to facilitate future policy query requests.
Interfaces	The interfaces for which the rule is defined. You can enter any combination of interface or interface role names, separated by commas. Enter the name or click **Select** to select the interface or interface role. **Note** If you leave the field blank, the query matches any interface or interface role.
Query for Global Rules	Whether the query should also consider global rules when querying access rules or inspection rules.
From Zone To Zone	For zone based firewall rules, the zones defined for the rule. Enter the zone names (which are interface roles), or click **Select** to select them from a list.
Actions	For zone based firewall rules, the actions defined for the rule.
Check if Matching Rules Are Shadowed by Rules Above	Whether to have the policy query results include rule conflict detection information. Selecting this option might have an impact on performance and cost results.

Interpreting Policy Query Results

Use the Policy Query Results dialog box to view the results of a policy query that you defined on the Query Device or Policy dialog box. The results report opens after you define your query parameters on the Querying Device or Policy Dialog Box, page 12-29 and click **OK**. For the procedure, see Generating Policy Query Reports, page 12-28. To see an example report, see Example Policy Query Result, page 12-34.

Tip In the query results table, you can double-click a row, or right-click and select **Go to Rule**, to select the rule in the rules policy page, where you can edit the rule. If the appropriate rules policy is not already selected in the policy selector, you might have to do this twice to actually select the rule.

To read the report, consider the following report sections:

- **Query Parameters**—The top portion of the report specifies the parameters you entered for the query. If you want to change them, click **Edit Query** to open the Querying Device or Policy Dialog Box, page 12-29, where you can make your changes and regenerate the report.

- **Results Table**—This table lists all rules that match your query. If you queried more than one type of rule, select the rule type you want to examine in the **Display** field. The columns in the table are the same as those for that type of rule, except for the following:
  - Match Status—Indicates how the rule matches your query:

    Complete Match—The rule matches all query parameters.

    Partial Match—All of the search criteria overlap or are a superset of the matched rule. For example, if you have a rule defined with a source address of 10.100.20.0/24, a destination address of 10.200.100.0/24 and a service of IP, and your query is to search for a source of 10.100.20.0/24, the match status is shown as a partial match because the query results represent only a portion of the rule's definition.

    No Effect—Rules are blocked by other matching rules, or a conflict exists that has no effect. For example, you might have two matching rules, A and B. If rule A's source address, destination address, and services are equivalent to, or contain, those of rule B, rule B is blocked by rule A. Thus, rule B will have no effect on traffic.

    In another example, you might have a global mandatory rule that permits a service, but a rule at the device (local) level denies the service. Because rules are recognized on a first-match basis, after discovering a match at the mandatory global scope, no other rules are checked. The local rule has no effect; the service is permitted, not denied. You should edit your policies to ensure you get the desired results.
  - Scope—Identifies whether a rule is shared or local, mandatory or default.
- **Details Table**—The details table shows the detailed query match information for the rule selected in the results table. The folders on the left represent the attributes for which you can see detailed information. Select a folder to view the details.

  The details show the query value, which is the parameter you defined, and the item in the rule that matches the parameter. The matching relationship is one of the following:
  - Identical—The parameter is identical to the value in the rule.
  - Contains—The parameter is a superset that contains the value in the rule. For example, the query parameter might have been a network/host object, and the rule used an IP address that was part of the object definition.
  - Is contained by—The parameter is a subset nested within the value of the rule.
  - Overlaps—The query parameter shows results that overlap between more than one policy object used in the rule. For example, the service query parameter was tcp/70-90 and the results show a service defined as tcp/80-100.

Related Topics

- AAA Rules Page, page 15-10
- Access Rules Page, page 16-10
- Inspection Rules Page, page 17-8
- Web Filter Rules Page (ASA/PIX/FWSM), page 18-3
- Zone-based Firewall Rules Page, page 21-58

Example Policy Query Result

Figure 12-4 shows an example of a policy query report on access rules. The criteria does not limit source, destination, service, and interface parameters, but limits the query to enabled rules. Both shared and local rules are included.

The Query Parameters section shows the query criteria for the report. In this example, the first row in the results table is selected, and the detailed information for that rule is shown in the details table in the bottom half of the window. In this example, the source folder is selected in the details table, and the result shows that the rule value, **any**, is an identical match to the query parameter *, which is equivalent to any source address.

For detailed information on reading the report, see Interpreting Policy Query Results, page 12-32.

Figure 12-4 Policy Query Results

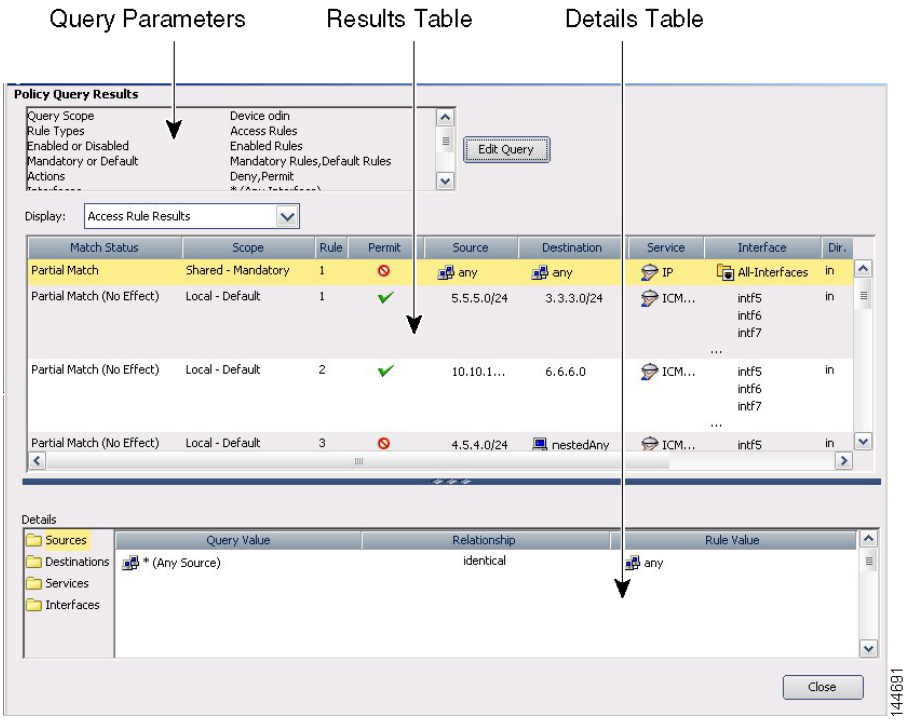

Related Topics

- Generating Policy Query Reports, page 12-28
- AAA Rules Page, page 15-10
- Access Rules Page, page 16-10
- Inspection Rules Page, page 17-8
- Web Filter Rules Page (ASA/PIX/FWSM), page 18-3
- Zone-based Firewall Rules Page, page 21-58

Optimizing Network Object Groups When Deploying Firewall Rules

When you deploy firewall rules policies to an ASA, PIX, FWSM, or IOS 12.4(20)T+ device, you can configure Security Manager to evaluate and optimize the network/host policy objects that you use in the rules when it creates the associated network object groups on the device. Optimization merges adjacent networks and removes redundant network entries. This reduces the runtime access list data structures and the size of the configuration, which can be beneficial to some FWSM and PIX devices that are memory-constrained.

For example, consider a network/host object named **test** that contains the following entries and that is used in an access rule:

```
192.168.1.0/24
192.168.1.23
10.1.1.0
10.1.1.1
10.1.1.2/31
```

If you enable optimization, when you deploy the policy, the resulting object group configuration is generated. Note that the description indicates the group was optimized:

```
object-group network test
 description (Optimized by CS-Manager)
 network-object 10.1.1.0 255.255.255.252
 network-object 192.168.1.0 255.255.255.0
```

If you do not enable optimization, the group configuration would be as follows:

```
object-group network test
 network-object 192.168.1.0 255.255.255.0
 network-object 192.168.1.23 255.255.255.255
 network-object 10.1.1.0 255.255.255.255
 network-object 10.1.1.1 255.255.255.255
 network-object 10.1.1.2 255.255.255.254
```

This optimization does not change the definition of the network/host object, nor does it create a new network/host policy object. If you rediscover policies on the device, the existing unchanged policy object is used.

Note If a network/host object contains another network/host object, the objects are not combined. Instead, each network/host object is optimized separately. Also, Security Manager cannot optimize network/host objects that use discontiguous subnet masks.

To configure optimization, select the **Optimize Network Object-Groups During Deployment** option on the Deployment Page, page 11-13 (select **Tools > Security Manager Administration** and select **Deployment** from the table of contents). The default is to not optimize network object groups during deployment.

Expanding Object Groups During Discovery

When you discover policies from a device that uses object groups, you can elect to have those object groups expanded into the items they contain rather than create policy objects from the group.

For example, if an object group named CSM_INLINE_55 contains the hosts 10.100.10.15, 10.100.10.18, and 10.100.10.25, importing an access control list by expanding the objects will create a rule that includes all three addresses in the source (or destination, as appropriate) cell rather than a network/host policy object named CSM_INLINE_55.

To configure expansion, you must have a naming scheme for your object groups that allows you to identify the prefix of groups that you want to expand. The default is to expand any object group that starts with the prefix CSM_INLINE. Configure these prefixes in the **Auto-Expand Object Groups with These Prefixes** field on the Discovery Page, page 11-25 by selecting **Tools > Security Manager Administration** and selecting **Discovery** from the table of contents.

Managing Identity-Aware Firewall Policies

Identity-aware firewall policies allow you to control traffic based on user identity or a host's fully-qualified domain name. For example, you can selectively allow a specific type of traffic for one user group while disallowing it for another user group, instead of allowing or disallowing all of that traffic. With fully-qualified domain names, you could disallow HTTP access to a specific server while allowing HTTP traffic to all other servers.

Identity awareness is integrated into several existing firewall rules; there is no unique identity-aware firewall policy. This chapter explains identity-aware firewall policies and how to implement them in the various policies that support identity awareness.

This chapter contains the following topics:

- Overview of Identity-Aware Firewall Policies, page 13-1
- Configuring Identity-Aware Firewall Policies, page 13-7
- Monitoring Identity Firewall Policies, page 13-27

Overview of Identity-Aware Firewall Policies

In traditional firewall policies, decisions are made based on source and destination IP addresses, ports, and services. The Identity Firewall in the ASA provides more granular control based on either or both of the following:

- User identity—You can configure access rules and security policies based on user names and user group names rather than through source IP addresses alone. The ASA applies the security policies based on an association of IP addresses to Windows Active Directory login information and reports events based on the mapped user names instead of network IP addresses.

  The Identity Firewall integrates with Microsoft Active Directory in conjunction with an external Active Directory (AD) agent that provides the actual identity mapping. The ASA uses Windows Active Directory as the source to retrieve the current user identity information for specific IP addresses and allows transparent authentication for Active Directory users. For information on setting up and configuring the AD agent, see *Installation and Setup Guide for the Active Directory Agent* on Cisco.com at http://www.cisco.com/en/US/products/ps6120/prod_installation_guides_list.html.
- FQDN network objects—You can use a host's fully-qualified domain name (FQDN) in a rule instead of its IP address. Thus, if the host's address changes (for example, because it acquires the address through DHCP), the rule will still apply.

Identity-based firewall services enhance the existing access control and security policy mechanisms by allowing users or groups to be specified as sources, and FQDNs in place of source or destination IP addresses. Identity-based security policies can be interleaved without restriction between traditional IP address based rules.

The key benefits of the Identity Firewall include:

- Decoupling network topology from security policies. The rules will apply to a user regardless of where the user connects in the network.
- Simplifying the creation of security policies.
- Providing the ability to easily identify user activities on network resources.
- Simplify user activity monitoring.

This section contains the following topics:

- User Identity Acquisition, page 13-2
- Requirements for Identity-Aware Firewall Policies, page 13-3
- Configuring the Firewall to Provide Identity-Aware Services, page 13-7

User Identity Acquisition

When you specify Active Directory user names or user group names in a firewall policy, the ASA eventually needs to map the name to an IP address to process packets. The ASA uses two primary sources for this information:

- User group membership—If you specify a user group in a rule, the ASA contacts the configured Active Directory (AD) server to obtain group membership.
- User-to-IP address mappings—For users who log into the network domain on your standard (non-VPN) network, the AD agent, in communication with the AD server, obtains the login information and creates a user-to-IP address mapping table. This information is supplied to the ASA.

You must install and configure the required AD servers and agents before you can configure user-based identity firewall policies. For an explanation of the various deployment scenarios, see the ASA configuration guide for ASDM or CLI at http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html.

User names are acquired for the following types of traffic, and include the AD domain unless noted otherwise:

- Standard traffic.
- Remote access VPN, including IPsec IKEv1 and IKEv2, AnyConnect clients, and L2TP VPN. If you use LDAP authentication for the VPN, and use the same server group for a domain for the VPN and identity firewall, the users are associated with the domain used for authentication. For all other authorization mechanisms, users acquired through VPN are considered to be in the LOCAL domain. The ASA reports these users to the AD agent, which distributes them to other ASAs or clients registered with the AD agent.

Note User names are not acquired for clientless SSL VPN.

- IPv4 cut-through proxy. User names are not acquired for IPv6 cut-through proxy. If the user includes the domain name during authentication, the user is associated with the domain name. Otherwise, the domain is the default domain as configured in the Identity Options policy. See Configuring Cut-Through Proxy, page 13-23.

Requirements for Identity-Aware Firewall Policies

Identity-aware firewall policies are not supported by all types of device and operating systems. The following table explains the requirements and some limits for implementing these types of policies in your network.

Table 13-1 Requirements for Identity-Aware Firewall Policies

Requirement	Description
Firewall device	ASA running ASA Software version 8.4(2) or higher, but not including the ASA-SM running 8.5(1). Single or multiple context configurations. **Tip** The ASA must have on-board encryption acceleration. To determine if the device has the required ability, log into the device console and use the **show version** command. The output should include "Encryption hardware device." You can register up to 100 ASAs with a single Active Directory agent.
Active Directory (AD)	You must use Active Directory to define users and user groups. The ASA obtains user group information directly from the AD server, which runs the LDAP protocol. You cannot use other types of LDAP servers. For detailed information on the types of AD servers supported, and their configuration requirements, see *Installation and Setup Guide for the Active Directory Agent* on Cisco.com at http://www.cisco.com/en/US/products/ps6120/prod_installation_guides_list.html. **Tip** You can have multiple AD servers, but they must have unique IP addresses among all domains. No other type of LDAP server is supported.

Table 13-1 Requirements for Identity-Aware Firewall Policies (continued)

Requirement	Description
AD agent	You must configure off-box AD agents to act as an intermediary between the ASA and the AD servers. The AD agent maintains an active mapping of users to IP address. By default, except for the 5505, the ASA obtains this list when it boots or reloads, and the AD agent sends new mappings as they are collected. The 5505 queries the AD agent on an as-needed basis in response to traffic matching rules that include identity criteria. We recommend that you use this default behavior, although you can change it using the Identity Options policy. The AD agent uses the RADIUS protocol. For information on setting up and configuring the AD agent, see *Installation and Setup Guide for the Active Directory Agent* on Cisco.com at http://www.cisco.com/en/US/products/ps6120/prod_installation_guides_list.html.
Client systems	Users who pass traffic through the device must use one of the following client platforms: • Windows XP SP3. • Windows Vista. • Windows 7. • Other systems that use Active Directory in a manner consistent with the explicitly supported platforms.
IPv6	IPv6 is supported with the following exceptions: • NetBIOS over IPv6 is not supported. • Multiple IPv6 addresses on user workstations is not supported. Windows 64-bit systems can use temporary IPv6 addresses when initiating some communications. If a user registers with the AD agent using one IPv6 address, then initiates communication with another address, an identity-aware firewall rule for the user would not be applied and instead a rule that matches the second IPv6 address would be applied. There are two options for disabling the use of these temporary addresses: – Disable IPv6 routing advertisements on all interfaces on all networking devices in the network. – On each Windows machine, open a command window, enter the following commands, and reboot the workstation: **netsh interface ipv6 set privacy state=disable** **netsh interface ipv6 set global randomizeidentifiers=disabled**

Table 13-1 Requirements for Identity-Aware Firewall Policies (continued)

Requirement	Description
NetBIOS logout probing (Optional.)	If you enable NetBIOS logout probing, the ASA can use NetBIOS to determine if an inactive user is logged off so the user can be removed from the database. The probe uses UDP-encapsulated NetBIOS traffic. Thus, you must ensure that access rules allow the following traffic on the networks between the ASA, AD agent, and user workstations: • Query packets: any UDP source port to UDP port 137 (UDP/137). • Query responses: UDP/137 source to any UDP port. In addition, you must configure workstations to provide user name information in NetBIOS reply packets. For Windows workstations, you need to enable the messenger service and configure WINS. If the messenger service is not turned on, the response from the workstation is the same whether the user is logged on or off. **Tips** • The NetBIOS logout probe is never used with VPN or cut-through proxy users. • The ASA has an inactive user timeout that is also used to remove users from the database. The timer applies to all user types. Thus, implementing NetBIOS probing is not required to remove inactive users from the database.

Table 13-1 Requirements for Identity-Aware Firewall Policies (continued)

Requirement	Description
DNS configuration (Required for fully-qualified domain name usage.)	If you use fully-qualified domain name (FQDN) network/host objects in firewall rules, you must configure the domain name system (DNS) settings as described in DNS Page, page 52-14. These settings identify the DNS servers used for looking up the names to determine the associated IP address. All processing is ultimately based on the IP address. When configuring DNS for FQDN usage, consider the following points: • DNS replies can be spoofed, which can open a security hole in your network. Specify only trusted DNS servers, ideally only those inside your network. • Some hosts can have constantly changing multiple IP addresses, so the ASA might not always have all valid IP addresses at any one time. • Host names with short time to live values will require frequent DNS lookups; this can impact the performance of the ASA. • Multiple host names can resolve to the same IP address. Ultimately, the firewall rules are applied based on IP address. Thus, if two names map to the same address, and your rules specify different services for those names, the service that is actually provided will be those specified on the first matched rule. Looked at another way, this means that you do not need to specify every version of an FQDN host name in your rules. When you know that several names always point to the same host, you can configure rules for the most commonly-used name and they will apply to all synonyms of that name.
Maximum limits	There are limits to the number of users, user groups, and IP addresses per user. If these limits are exceeded, identity-aware processing will not occur for the additional traffic: • IP address limits—A user can be associated with at most 8 IP addresses across all domains. • User group limits—Policies can be applied to up to 256 user groups. Users can be in multiple user groups. • User limits—Policies can be applied to up to the following number of users. This number is the total aggregate across all contexts defined on the device. – ASA 5505—1024 users. – Other ASA 5500 series—64,000 users.

Configuring the Firewall to Provide Identity-Aware Services

To provide identity-aware firewall services to your network, you need to configure several policies to enable the firewall to process user-based or fully-qualified domain name (FQDN)-based rules. The ASA depends on other servers in your network to provide the user, user group, and FQDN name resolution services required to implement your identity-aware policies.

The required configuration depends on which aspects of identity awareness that you will use:

- User, user group resolution—To use identity user group objects in your firewall rules, you must configure several objects and policies to identify the Active Directory servers that will supply user and user group information.
- FQDN resolution—To use FQDN network/host objects in your firewall rules, you must configure DNS servers to resolve FQDNs to IP addresses.

This procedure explains the overall process for implementing identity-aware policies.

Before You Begin

Your network must meet the requirements explained in Requirements for Identity-Aware Firewall Policies, page 13-3. The following procedure assumes that you are already using Active Directory (AD), that you have installed and configured the AD agents, and that these services are working correctly.

Step 1 Enable AD user and user group resolution.

a. Create the policy objects needed to identify the AD servers and agents and configure the NetBIOS domain for the server groups. For detailed information, see Identifying Active Directory Servers and Agents, page 13-8.

b. If you want non-default settings, change the identity options. Use these options to enable the NetBIOS logout probe and to configure various timers and error handling. For detailed information, see Configuring Identity Options, page 13-15.

c. If you want to create user groups defined on the ASA (in addition to AD-defined user groups), create the required identity user group policy objects. See Creating Identity User Group Objects, page 13-19.

Step 2 Enable FQDN network/host object resolution.

a. Configure DNS servers in the DefaultDNS group. DNS is required to resolve FQDNs to IP addresses. For information on configuring DNS, see DNS Page, page 52-14.

b. Create FQDN network/host objects as described in Creating Networks/Hosts Objects, page 6-82.

Step 3 Configure firewall rules to use FQDN objects, usernames, user group names, or identity user group objects. See Configuring Identity-Based Firewall Rules, page 13-21.

Step 4 Monitor the identity firewall system. See Monitoring Identity Firewall Policies, page 13-27.

Configuring Identity-Aware Firewall Policies

Identity awareness is integrated into several existing firewall rules; there is no unique identity-aware firewall policy. The topics in this section explain the various procedures for integrating identity awareness into firewall policies.

This section contains the following topics:

- Enabling Identity-Aware Firewall Services, page 13-8
- Creating Identity User Group Objects, page 13-19
- Selecting Identity Users in Policies, page 13-21
- Configuring Identity-Based Firewall Rules, page 13-21
- Configuring Cut-Through Proxy, page 13-23
- Collecting User Statistics, page 13-25
- Filtering VPN Traffic with Identity-Based Rules, page 13-26

Enabling Identity-Aware Firewall Services

Use the Identity Options policy to enable identity-aware firewall services. To configure the policy, do one of the following:

- (Device view) Select an ASA device, then select **Identity Options** from the Policy selector.
- (Policy view) Select **Identity Options (ASA)** from the Policy selector. Select an existing policy or create a new one.

The policy includes the following tabs:

- **AD Setup**—Configure the Active Directory servers that define the users and user groups for the network, and the AD agent used to collect the information and provide it to the ASA. See Identifying Active Directory Servers and Agents, page 13-8.
- **Advanced**—Enable or disable user identity services and configure options for error handling, the NetBIOS logout probe, idle timeout, and AD agent communication settings. See Configuring Identity Options, page 13-15.

Identifying Active Directory Servers and Agents

Use the AD Setup tab of the Identity Options policy to identify the Active Directory servers and agents to use for user identity information. You must configure at least one AD server and one AD agent to enable identity-aware firewall policies that include user specifications (such as identity user group objects).

Note ASA Software 8.4(2+) is required for identity-aware firewall.

Before You Begin

The configuration uses AAA server group policy objects, and the server group objects incorporate AAA server objects. You can create these objects through the Policy Object Manager (by selecting Manage > Policy Objects), or you can create them while completing this procedure (by using the configuration wizard or by clicking the Add Object + button in the object selector dialog box).

The objects need to meet these requirements:

- AD servers—Must use the LDAP protocol. If you select Microsoft as the LDAP server type, you can also specify the LDAP Group Base DN to identify the base directory for user group searches, to reduce search time. If you select Auto Detect, you cannot configure the group base DN, even though

Microsoft AD servers are the only type of LDAP server that you can use in the identity firewall configuration. You must also abide by the following limitations for communications between Security Manager and Active Directory:

  - Do not select the Enable LDAP over SSL option.
  - Do not select the SASL Kerberos Authentication option. Only simple and SASL MD5 authentication mechanisms are supported. The simple mechanism, in which usernames and passwords are transmitted in clear text, is used if you do not select one of the SASL options.

- AD agents—Must use the RADIUS protocol. In the AAA server group object, select the **AD Agent Mode** option.

  You should install the AD agents and configure them prior to configuring this policy. You can configure at most two AD agents in the server group: the second agent is used only if the first agent ceases to respond to queries. Any agents defined after the first two are ignored.

  Obtain the AD agent software from http://www.cisco.com/go/asa. For information on setting up and configuring the AD agent, see *Installation and Setup Guide for the Active Directory Agent* on Cisco.com.

Related Topics

- Requirements for Identity-Aware Firewall Policies, page 13-3
- Understanding AAA Server and Server Group Objects, page 6-27
- Creating AAA Server Objects, page 6-32
- AAA Server Dialog Box—LDAP Settings, page 6-40
- Creating AAA Server Group Objects, page 6-48
- Configuring Identity Options, page 13-15

Step 1 Do one of the following:

- (Device view) Select an ASA device, then select **Identity Options** from the Policy selector. Select the **AD Setup** tab.
- (Policy view) Select **Identity Options (ASA)** from the Policy selector. Select an existing policy or create a new one. Select the **AD Setup** tab.

Step 2 If you want to be guided through the AD setup, click the **Configure Identity** button to start the Identity configuration wizard. The wizard walks you through the process of configuring the AD servers for a domain, and the AD agents, and can create the required AAA server and AAA server group objects for you.

The wizard goes through the following steps:

- AD Server Settings—To configure the AD servers for a domain. See Identity Configuration Wizard Active Directory Settings, page 13-11.
- AD Agent Settings—To configure the AD agents for the ASA. See Identity Configuration Wizard Active Directory Agent, page 13-13.
- Preview—To show you which objects will be created. See Identity Configuration Wizard Preview, page 13-15.

Tip You can use the wizard multiple times to configure different NetBIOS domains. However, the wizard always prompts for AD agent information. Because you can configure a single group for AD agents, not a separate group per domain, the selection overwrites any AD agent configuration that you have already made. So be sure to select the same AAA server group for the AD agents each time you run the wizard.

Step 3 If not using the wizard, configure the AD servers. The AD servers are used for obtaining user membership information for any AD user groups that you use in identity-aware firewall policies.

The table lists the AD servers for the network. You need to add an entry for each NetBIOS domain name. Each row defines the AAA server group used to identify the AD LDAP servers for the domain, and whether identity-aware firewall rules for the domain are active or inactive if the AD server group is unavailable.

You can do the following:

- To add an entry, click the **Add Row (+)** button and fill in the Add AD Domain Server dialog box. See Domain AD Server Dialog Box, page 13-10.
- To edit an entry, select it and click the **Edit Row (pencil)** button.
- To delete an entry, select it and click the **Delete Row (trash can)** button.

Step 4 If not using the wizard, configure the AD agents. The AD agents obtain user login/logoff and IP address mappings from the AD servers. The ASA then obtains the information from the AD agent.

In **Active Directory Agent Group**, enter the name of the AAA server group object that defines the list of AD agents, or click **Select** to select it from a list or to create a new group object.

Step 5 In **Default Domain**, select the domain to configure as the default domain on the device. You must add the domain to the AD server table before you can select it as the default domain.

The default is LOCAL, which applies to user groups defined on the device or to VPN users who authenticate using a method other than an AD server configured for identity services. This setting is also used if you configure cut-through proxy (see Configuring Cut-Through Proxy, page 13-23).

Step 6 Click **Save** to save your changes.

You are asked if the identity settings page in the administrative settings should be updated with the domain-to-AD server mappings. The identity settings determine which servers are used when you use the Find feature when specifying users or user groups in a firewall policy or an identity user group object. The identity administrative settings do not affect the configuration of the ASA.

Domain AD Server Dialog Box

Use the Add or Edit Domain AD Server dialog box to define the Active Directory server group for a NetBIOS domain. If you configure firewall rules for a user group in the NetBIOS domain, the user membership is determined by querying the AD servers defined for the domain.

Navigation Path

Do one of the following:

- From the AD Setup tab of the Identity Options page, click the Add or Edit buttons for the domain table. See Identifying Active Directory Servers and Agents, page 13-8.

- From the Identity Settings Security Manager Administration page, click the Add or Edit buttons for the settings table. These settings determine which servers are used when using Find to locate a user or user group name when configuring firewall rules or identity user group objects. See Identity Settings Page, page 11-38.

Field Reference

Table 13-2 Domain AD Server Dialog Box

Element	Description
Domain	The NetBIOS domain for this AD server group. The domain name can be up to 32 characters, typically in all uppercase. For example, if the user specification is DOMAIN\user1, DOMAIN is the NetBIOS domain name.
AD Server Group	The name of the AAA server group policy object that identifies the AD servers for this domain. The object must use the LDAP protocol. Click **Select** to select the object or to create a new one.
Disable Rules When Server Is Down (Identity Options policy only.)	Whether to disable all identity-aware firewall rules for this domain if the domain controller is down. If you select this option, all users for a domain are marked as disabled until the server becomes available.
Update Administrative Settings (Identity Options policy only.)	Whether to add the domain and server mapping to the Security Manager Administration Identity Settings page. This administrative page determines which AD servers are queried when you try to find users or user groups when adding them to firewall policies or to identity user group objects. For more information, see Identity Settings Page, page 11-38.

Identity Configuration Wizard Active Directory Settings

Use the Active Directory Settings page of the Identity Configuration wizard to identify the Active Directory (AD) servers for a NetBIOS domain. These settings are required to enable user-identity-aware firewall policies for users in the domain.

Navigation Path

Do one of the following:

- From the AD Setup tab of the Identity Options page, click the **Configure Identity** button. See Identifying Active Directory Servers and Agents, page 13-8.
- If the Identity Options policy is not already configured, you can start the wizard from the AAA Rules, Access Rules, or Inspection Rules policies by clicking the Select button for the User field and then clicking Yes when asked if you want to configure identity.

Field Reference

Table 13-3 Identity Configuration Wizard Active Directory Settings

Element	Description
NetBIOS Domain	The NetBIOS domain for this AD server group. The domain name can be up to 32 characters, typically in all uppercase. For example, if the user specification is DOMAIN\user1, DOMAIN is the NetBIOS domain name.
Select Existing AD Server Group	Select this option if the AAA server group policy object that identifies the required AD servers already exists. The object must use the LDAP protocol. Click **Select** next to the Group Name field to select the object.
Create New AD Server Group	Select this option if the AAA server group policy object does not already exist, or you want the wizard to create a new object. Configure the remaining options to identify the group and the servers that it contains.
Create AD Server Group Properties	
Group Name (When creating the group in the wizard.)	The name of the AAA server group object that you want to create. The name can be up to 16 characters.
AD Server Name/IP	One of the following: • The name of an existing AAA server object that defines the AD server. Click **Select** to select the object from a list. If you select an object, you cannot configure the remaining properties. • The IP address of the AD server.
Username	The name of the user or the directory object in the LDAP hierarchy used for authenticated binding (maximum of 128 characters). Authenticated binding is required by some LDAP servers (including the Microsoft Active Directory server) before other LDAP operations can be performed. This field describes the authentication characteristics of the device. These characteristics should correspond to those of a user with administrator privileges. This string is case-sensitive. Spaces are not permitted in the string, but other special characters are allowed. Typically, this is a username such as DOMAIN\Administrator. However, you can use the more traditional format too, for example, cn=Administrator,OU=Employees,DN=example,DN=com.
Password Confirm	The case-sensitive, alphanumeric password for accessing the LDAP server (maximum of 64 characters). Spaces are not allowed.

Table 13-3 Identity Configuration Wizard Active Directory Settings (continued)

Element	Description
Interface	The interface whose IP address should be used for all outgoing packets (known as the source interface). Enter the name of an interface or interface role, or click **Select** to select it from a list or to create a new interface role. **Tips** • If you enter the name of an interface, make sure the policy that uses this AAA object is assigned to a device containing an interface with this name. • If you enter the name of an interface role, make sure the role represents a single interface, not multiple interfaces. • Only one source interface can be defined for the AAA servers in a AAA server group, so if you specify more than one server, ensure that they all use the same interface.
Add Another AD Server	Click this button only if you want to create an additional server. When you click the button, the information in the server fields is saved and the fields are cleared so that you can add information about the next server. You can add up to 16 servers in single-context mode and 4 servers in multiple-context mode.

Identity Configuration Wizard Active Directory Agent

Use the Active Directory Agent Settings page of the Identity Configuration wizard to identify the Active Directory (AD) agents for a NetBIOS domain. These settings are required to enable user-identity-aware firewall policies for users in the domain.

Tip You can configure a single AD agent group for an ASA; you do not configure a different group for each NetBIOS domain. Thus, if you already configured the correct AD agent group in the Identity Options policy, select the same group on this wizard page. Your selection here will replace the group defined in the policy.

Navigation Path

Do one of the following:

- From the AD Setup tab of the Identity Options page, click the **Configure Identity** button and proceed to this page. See Identifying Active Directory Servers and Agents, page 13-8.
- If the Identity Options policy is not already configured, you can start the wizard from the AAA Rules, Access Rules, or Inspection Rules policies by clicking the Select button for the User field and then clicking Yes when asked if you want to configure identity.

Field Reference

Table 13-4 Identity Configuration Wizard Active Directory Agent Settings

Element	Description
Select Existing AD Agent Group	Select this option if the AAA server group policy object that identifies the required AD agents already exists. The object must use the RADIUS protocol, and should have the option **AD Agent Mode** selected. Click **Select** next to the Group Name field to select the object.
Create New AD Agent Group	Select this option if the AAA server group policy object does not already exist, or you want the wizard to create a new object. Configure the remaining options to identify the group and the servers that it contains.
Create AD Agent Group Properties	
Group Name (When creating the group in the wizard.)	The name of the AAA server group object that you want to create. The name can be up to 16 characters.
AD Agent Name/IP	One of the following: • The name of an existing AAA server object that defines the AD agent. Click **Select** to select the object from a list. If you select an object, you cannot configure the remaining properties. • The IP address of the AD agent.
Secret Key Confirm	The shared secret that is used to encrypt data between the network device (client) and AAA server. The key is a case-sensitive, alphanumeric string of up to 127 characters. Special characters are permitted. The key you define in this field must match the key on the RADIUS server. Enter the key again in the Confirm field. If you do not define a key, all traffic between the AAA server and its AAA clients is sent unencrypted.
Interface	The interface whose IP address should be used for all outgoing packets (known as the source interface). Enter the name of an interface or interface role, or click **Select** to select it from a list or to create a new interface role. **Tips** • If you enter the name of an interface, make sure the policy that uses this AAA object is assigned to a device containing an interface with this name. • If you enter the name of an interface role, make sure the role represents a single interface, not multiple interfaces. • Only one source interface can be defined for the AAA servers in a AAA server group, so if you specify more than one server, ensure that they all use the same interface.

Table 13-4 Identity Configuration Wizard Active Directory Agent Settings (continued)

Element	Description
Add Secondary AD Agent	Click this button only if you want to create an additional agent. The agent is used in case the first agent becomes unavailable. When you click the button, the information in the agent fields is saved and added to the preview pane, and the fields are cleared so that you can add information about the secondary agent.

Identity Configuration Wizard Preview

Use the Preview page of the Identity Configuration wizard to verify the information you entered into the wizard.

The preview summarizes the objects that will be created or used for the Active Directory configuration for the NetBIOS domain.

- AD server group shows the name of the AAA server group object for the AD servers used in the domain. The table shows the AAA server objects that define each of the AD servers.
- AD Agent shows the name of the AAA server group object for the AD agents. The primary and secondary agent shows the AAA server object that defines the agents.

For objects that the wizard will create, names are automatically generated for the AAA server objects, either adding **ldap_** or **radius_** as a prefix to the server IP address.

To make changes, click **Back**. Otherwise, click **Finish** to save the settings.

Tip After you complete the wizard, you can edit the properties of the newly-created objects to configure settings that the wizard left as default settings.

Navigation Path

Do one of the following:

- From the AD Setup tab of the Identity Options page, click the **Configure Identity** button and proceed to this page. See Identifying Active Directory Servers and Agents, page 13-8.
- If the Identity Options policy is not already configured, you can start the wizard from the AAA Rules, Access Rules, or Inspection Rules policies by clicking the Select button for the User field and then clicking Yes when asked if you want to configure identity.

Configuring Identity Options

Use the Advanced tab of the Identity Options policy to enable or disable user identity services and configure options for error handling, the NetBIOS logout probe, idle timeout, and AD agent communication settings. The options on this tab have default values, so you need to change them only if you want to fine-tune the settings for your network.

Navigation Path

- (Device view) Select an ASA device, then select **Identity Options** from the Policy selector. Select the **Advanced** tab.
- (Policy view) Select **Identity Options (ASA)** from the Policy selector. Select an existing policy or create a new one. Select the **Advanced** tab.

Related Topics

- Identifying Active Directory Servers and Agents, page 13-8
- Requirements for Identity-Aware Firewall Policies, page 13-3

Field Reference

Table 13-5 Identity Options Advanced Tab

Element	Description
Enable User Identity	Whether to enable the device to obtain user identity information from the AD agent and AD servers, if they are configured on the **AD Setup** tab. The default is enabled. If you change this option and deploy, the change has the following effect based on the new setting: • Disabled—The entire IP address to user mapping database is flushed and all users without activated user-specific rules are released. The AD agent and servers are no longer queried for updates, and all activated user-identity-based rules will have no effect on traffic. • Enabled—Activated users are recreated gradually through communications with the AD agent. VPN users might need to reauthenticate. Queries to the AD agent and AD server recommence.
Error Conditions	
Disable Rules When Active Directory Agent Is Down	Whether to disable all rules that include user identity if the connection to the AD agent is unavailable. If you select this option, all user-to-IP address mappings are marked disabled and all rules that include user specifications are not applied to traffic. By default the option is disabled.
Remove User IP When NetBIOS Probe Fails	Whether to remove the User's IP address mapping from the database if the NetBIOS probe for the user fails for any reason, whether the probe is somehow blocked in the network or the probe fails because the user is not in operation. The user must log into the workstation again. This option has effect only if you enable the NetBIOS logout probe on this page. By default the option is disabled.
Remove User IP When User's MAC Address is Inconsistent	Whether to check the Media Access Control (MAC) address in each request from a user-mapped IP address to the MAC address in the previous packet. If you select this option, and the MAC address changes between packets, the user-to-IP address mapping is removed from the database, subsequent packets are dropped, and the user must reauthenticate to Active Directory. The AD agent is notified if the user-to-IP mapping is removed due to MAC mismatch. By default this option is enabled. MAC checking occurs only on packets from IP addresses on networks that are directly attached to the ASA. VPN users are not checked.
Track User Not Found	Whether to enable user-not-found tracking. By default, the option is disabled.

Table 13-5 Identity Options Advanced Tab (continued)

Element	Description
NetBIOS Logout Probe	
Enable (NetBIOS Logout Probe)	Whether to enable the NetBIOS logout probe. You can use the probe to proactively determine if a user has logged out of the network, allowing the device to remove the user-to-IP address mapping more quickly than if idle timeout is the only mechanism used for this purpose. By default the probe is disabled, and users are removed only if they are idle for longer than the Idle Timeout value. Users are probed only if they are in the active state and they are used in at least one activated rule. VPN and cut-through proxy users are not probed. The AD agent is notified if the user-to-IP mapping is removed by the NetBIOS logout probe. In addition to configuring the following options, see Requirements for Identity-Aware Firewall Policies, page 13-3.
Probe Timer	The frequency of sending NetBIOS probes to activated users, regardless of whether the user is idle. The default is 15 minutes, the range is 1 to 65535 minutes.
Retry Interval	The frequency of retrying the probe if a response is not received from an IP address, and the number of times the probe should be retried. The default is 3 seconds and 3 retries. The range is 1 to 65535 seconds, for retry count, 1 to 256. If there is no response from the final retry, the user-to-IP address mapping is removed if you selected the **Remove User IP When NetBIOS Probe Fails** option; otherwise, the address is probed during the next interval.
User Name	When a NetBIOS response is received, how to handle the response based on the usernames returned: • **Match Any** (the default)—Any username in the response can match the username in the database for the IP address. If there are multiple names in the response (that is, more than one user is logged into the workstation), if any user in the response matches a user in the database, the probe is considered successful and the mapping is retained. • **User Not Needed**—The usernames in the NetBIOS response are ignored; the query response is sufficient to maintain the user-to-IP address mapping. This option is useful if the messenger service is not turned on in the workstation, in which case the NetBIOS response will not contain usernames. The option is also useful when multiple users log into a workstation. • **Exact Match**—There must be one username in the NetBIOS response, and it must exactly match the username in the user-to-IP address mapping database. If there is more than one user, or if the username does not match, the mapping is removed from the database and the IP address is marked as inactive.

Table 13-5 Identity Options Advanced Tab (continued)

Element	Description
Users	
Idle Timeout	The amount of time, in minutes, to allow the user to be idle before removing the user-to-IP address mapping in the database. Once removed, the user must log in again to update the mapping (for example, by using Ctrl+Alt+Delete to lock the workstation, then log in again). The default is 60 minutes, the range is 1 to 65535 minutes. You can deselect the option to disable idle timeout checking, in which case user-to-IP mappings are not removed due to idleness. VPN and cut-through proxy users are not subject to this timer. The AD agent is not notified if the user-to-IP address mapping is removed due to idle timeout.
Active Directory Agent	
Hello Timer	The frequency of sending Hello packets to the AD agent. The ASA uses Hello packets to obtain ASA replication status and domain status. If the ASA does not receive a response after the final retry, the AD agent is considered down, and the ASA switches to the backup AD agent, if you configure one. By default, Hello packets are sent every 30 seconds, and up to 5 retries are attempted if no response is received. The range is 10 to 65535 seconds and 1 to 65535 retries.
Poll Groups Timer	How often the Active Directory server should be queried to obtain user membership lists for user groups that you have specified in firewall rules. The ASA queries the server for membership in a group only if you have used the group; it does not query every group defined in the AD server. The default is 8 hours, the range is 1 to 65535 hours. **Tip** If group membership changes, the changes are not reflected in rule processing until this timer expires and the ASA polls the AD server for updated information. Thus, you should configure the timer based on the frequency of changes to group membership in your network, balancing the need to update group membership in the ASA with the desire to reduce the amount of polling.

Table 13-5 Identity Options Advanced Tab (continued)

Element	Description
Retrieve User Information	How the ASA should retrieve user-to-IP address mappings from the AD agent. • **Full Download** (default for ASA non-5505 devices)—On boot, the ASA obtains the full user-to-IP address mapping database from the AD agent, and then gets incremental updates as users log into and out of the network. Use this option on the 5505 only if there are fewer than 1024 users in the network, because the 5505 is limited to 1024 user-to-IP mappings. For the 5505, the default On Demand setting is appropriate if only a few users will pass traffic through the device. • **On Demand** (default for ASA 5505 devices)—The ASA queries the AD agent for user-to-IP mappings only when a new packet requires a connection and no mapping exists. This option uses less memory, but there can be a delay in getting the mapping, and the packets are initially evaluated based on traditional source and destination IP address and service information, which might result in the wrong action. The potential delay can be increased if a large number of users log in at the same time, either due to corporate culture or to a malicious attack.

Creating Identity User Group Objects

You can create identity user group objects to identify individual users, user groups, or a combination of users and groups. These users and groups must be defined in Active Directory (AD), you cannot define other types of users.

Tip Identity user groups are defined on the ASA. You do not need to create these groups to duplicate groups that are already defined in AD. You can directly specify AD groups in firewall rules. Identity user group objects are needed only to define collections of users and user groups that do not otherwise exist in AD.

There are two pre-defined identity user groups. These groups are used when configuring cut-through proxy, as described in Configuring Cut-Through Proxy, page 13-23.

- all-auth-users—To match any IP address that has been associated with an authenticated user.
- all-unauth-users—To match only IP addresses that have **not** been associated with authenticated users.

Tips

- Use of these objects is supported on ASA 8.4(2+) only.
- You must configure the Identity Options policy on the ASA to enable the use of these objects.
- You can create identity user group objects when defining policies or objects that use this object type. For more information, see Selecting Identity Users in Policies, page 13-21.

Related Topics

- Configuring Identity-Based Firewall Rules, page 13-21

- Requirements for Identity-Aware Firewall Policies, page 13-3
- Identity Settings Page, page 11-38
- Creating Policy Objects, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **Identity User Group** from the Object Type selector.

Step 3 Right-click in the work area, then select **New Object** to open the Identity User Group dialog box.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Add and remove items in the **Members in Group** list to identify the users and user groups defined in the object.

To populate the list, do any combination of the following:

- In **Available Identity User Group**, select an existing object and click the **Add >>** button between the lists.
- In **Search User/User Group**, select a user or user group from the Active Directory server configured for the domain in the Identity Settings administration options. You must configure the settings before you can select users or user groups, so that Security Manager knows which AD server to use.

  To find a user or user group, select the NetBIOS domain, indicate whether you are searching for a user or user group, and enter a search string. Then, click **Search** to find matches. A name is considered a match if the string appears anywhere within the name (first, middle initial, last), user ID, CN, or for groups, user group name.

  To add the user or group, select it in the list and click the **Add >>** button between the lists.
- In **Type in comma separated identity user or user group**, type in a valid name, then click the **Add >>** button between the lists. Separate multiple names with commas; they are added as separate lines in the members list.

  You can enter names in the following formats:

  – Individual users: NETBIOS_DOMAIN\user

  – User groups (note the double \): NETBIOS_DOMAIN\\user_group

  If you do not include the domain name, one is added for you based on the options selected in the Security Manager Administration Identity Settings page. If you precede the name with \ or \\, the default domain defined on the Identity Settings page is automatically added.
- To remove an item from the object, select it in the Members list and click the **<< Remove** button between the lists.

Step 6 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 7 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 8 Click **OK** to save the object.

Selecting Identity Users in Policies

In any policy or policy object that allows the specification of identity users, whether directly or through the selection of an identity user group object, you can click the Select button next to the User field to help you enter the information.

In the Identity User Group Selector dialog box, you can define the content of the User field by populating the **Members in Group** list. To populate the list, do any combination of the following:

- In **Available Identity User Group**, select an existing object and click the **Add >>** button between the lists. If the desired object does not exist, you can click the **Add (+)** button below the list to create a new object. You can also select an object and click the **Edit (pencil)** button to modify it or to examine its contents.

  There are two pre-defined identity user groups. These groups are used when configuring cut-through proxy, as described in Configuring Cut-Through Proxy, page 13-23.

  - all-auth-users—To match any IP address that has been associated with an authenticated user.
  - all-unauth-users—To match only IP addresses that have **not** been associated with authenticated users.

- In **Search User/User Group**, select a user or user group from the Active Directory server configured for the domain in the Identity Settings administrative options. You must configure the settings before you can select users or user groups, so that Security Manager knows which AD server to use.

  To find a user or user group, select the NetBIOS domain, indicate whether you are searching for a user or user group, and enter a search string. Then, click **Search** to find matches. A name is considered a match if the string appears anywhere within the name (first, middle initial, last), user ID, CN, or for groups, user group name.

  To add the user or group, select it in the list and click the **Add >>** button between the lists.

- In **Type in comma separated identity user or user group**, type in a valid name, then click the **Add >>** button between the lists. Separate multiple names with commas; they are added as separate lines in the members list.

  You can enter names in the following formats:

  - Individual users: NETBIOS_DOMAIN\user
  - User groups (note the double \): NETBIOS_DOMAIN\\user_group

  If you do not include the domain name, one is added for you based on the options selected in the Security Manager Administration Identity Settings page as explained in Identity Settings Page, page 11-38. If you precede the name with \ or \\, the default domain defined on the Identity Settings page is automatically added.

- To remove an item from the object, select it in the Members list and click the **<< Remove** button between the lists.

Configuring Identity-Based Firewall Rules

Identity awareness is integrated into the access control entries, or rules, in the ACLs used to provide firewall services. Because the feature is integrated into the ACL, the techniques for adding identity-based rules to a firewall policy are the same for all types of firewall policy. This topic provides general guidance on how to incorporate identity-based rules into your existing policies, and directs you to more specific information on configuring each type of policy that allows identity-based rules.

Guidelines For Adding Identity-Based Rules

Following are some general guidelines and recommendations for adding identity-based rules:

- FQDN (fully-qualified domain name) network/host objects are allowed in both Source and Destination fields. For information on configuring these objects, see Creating Networks/Hosts Objects, page 6-82.
- User, user group, and identity user group objects, which specify Active Directory (AD) user or user group names, are defined in a separate field: User. If you configure a rule with one or more user names, user group names, or identity user group objects, the specifications modify the Source address configuration only. They never apply to the addresses specified in the Destination field. For information on configuring these identity user group objects, see Creating Identity User Group Objects, page 13-19.

  You must always configure a source address in a rule, even if you want the rule to primarily operate based on the specified users or user groups. The source and user specifications conjoin to control the scope of the rule. Based on the value of the Source field, the rules operate as follows:

  - **Source = any**—Use "any" as the source if you want the rule to apply based solely on the user specifications. These rules will match the user specification regardless of the workstation IP address from which the user sends traffic.
  - **Source = anything else**—If you specify anything other than "any" as the source address, the rule applies only if the user sends traffic from an IP address that matches the source address specification. Use this technique if you want to provide variable services based on the source network.

    For example, you might have an internal trusted network from which you would allow access to a sensitive destination for users in a particular user group, although you would deny access even to those users if they were outside of the trusted network. In this case, you would create a permit rule that specified the trusted network as the source, the trusted user group as the user, and the sensitive server as the destination. You could also create a specific deny rule with just the source and destination specified, or allow the default deny any rule to capture the non-matching traffic.
- Evaluate whether there are classes of traffic that will never be sensitive to user identity. For example, you might allow DNS traffic for all users. Place these types of rules above identity-based rules so that matching traffic can be quickly allowed before the device needs to evaluate identity-based rules.
- When troubleshooting rules, keep in mind that ultimately rules are applied based on IP address. FQDN rule matching is based on DNS lookups, and the IP address of a host can change between a successful lookup and the next time the lookup is refreshed. For users, IP address mappings are obtained from the AD agent configured in the network or by authentications conducted by the ASA itself.
- FQDN and user specifications are completely independent. You can use one without the other.

Firewall Policies That Allow Identity-Based Rules

Identity-based rules are allowed on ASA 8.4.2+ only. The following policies allow you to configure identity-based rules:

- AAA Rules—Select **Firewall > AAA Rules** and see Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4.

Tip You can use AAA rules to configure cut-through proxy, which allows users whose IP address mappings have become invalid, resulting in denied network access, to authenticate directly to the ASA to resolve the mapping problem. See Configuring Cut-Through Proxy, page 13-23.

- Access Rules—Select **Firewall > Access Rules** and see Configuring Access Rules, page 16-7.
- Inspection Rules—Select **Firewall > Inspection Rules** and see Configuring Inspection Rules, page 17-5.
- Policies that use extended ACL policy objects—Several firewall policies use extended ACL policy objects to define traffic matching criteria instead of incorporating a rule table directly in the policy. You can configure extended ACL policy objects to include FQDN objects or user specifications (see Creating Extended Access Control List Objects, page 6-54). You can then use these identity-based extended ACL objects in the following policies:
  - Botnet Traffic Filter Rules—Select **Firewall > Botnet Traffic Filter Rules** and see Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6. You can use identity-based ACLs as traffic classification for Enable and Drop rules.
  - IPS, QoS, and Connection Rules (service policy rules)—Select **Platform > Service Policy Rules > IPS, QoS, and Connection Rules** and see Service Policy Rules Page, page 58-5.

    Traffic match criteria in this policy is based on extended ACL policy objects that are incorporated into traffic flow policy objects. You must select one of the options for specifying an ACL in the traffic flow object to incorporate identity-based traffic classification. You can use identity-based ACLs for all service types. For more information, see Configuring Traffic Flow Objects, page 58-18.

    One of the services available in this policy, User Statistics, is specifically designed to collect accounting information for identity-based firewall users. See Collecting User Statistics, page 13-25.
  - VPN filter in remote access group policies—The VPN filter ACL is applied to VPN traffic. You can configure a VPN filter on the Connection Settings page in an ASA Group Policy object, which you use in a remote access connection policy. See ASA Group Policies Connection Settings, page 34-33 and Filtering VPN Traffic with Identity-Based Rules, page 13-26.

Policies That Do Not Allow Identity-Based Rules or Objects

There are several types of policy where you can specify network/host objects or extended ACL objects, but where the policy does not allow FQDN network/host objects or ACLs that use those objects or identity user group objects. Following are some examples where you cannot use these types of objects:

- Routing policies, including route maps.
- Network address translation (NAT).
- WCCP (web cache control protocol).
- Crypto maps in VPN configurations.
- Dynamic access policies in remote access VPN configurations.

Configuring Cut-Through Proxy

When you use identity-aware firewall policies, user-to-IP address mappings are obtained from various facilities, primarily from the AD agent in the network. Although mappings are updated regularly, there can occur instances where a firewall rule blocks a legitimate user because the user-to-IP address mapping is not synchronized.

You can configure cut-through proxy to account for this possibility. With cut-through proxy, if a user is blocked, the user can sign on directly to the ASA, and the ASA will update the user-to-IP mapping to correctly reflect the current IP address for the user. The new mapping is forwarded to all contexts that contain the interface where the HTTP/HTTPS packets are received and authenticated.

You use AAA rules to configure cut-through proxy. You have two configuration choices, based on whether there is one or more NetBIOS domains in the network:

- Single domain—Configure a regular AAA rule for authentication and specify the LDAP server group that identifies the Active Directory servers for the domain. Use "any" for the source, and the IP address of the ASA for the destination. For service, you can include HTTP and HTTPS. Then, when the user needs to authenticate to the server, the user enters one of the standard authentication URLs, where *interface_ip* is the IP address of the interface and *port* is optionally the port number, if you specify a non-default port for the protocol in the interactive authentication table: **http://***interface_ip*[**:***port*]**/netaccess/connstatus.html** or **https://***interface_ip*[**:***port*]**/netaccess/connstatus.html**.

Tip The user-to-IP mapping is put under the same domain as configured for the selected AD server group. If you use another means for authentication, the mapping is placed under the LOCAL domain.

- Multiple domains—Configure two authentication rules that use the User-Identity option instead of a specific AAA server group. The following procedure explains this setup. Note that this setup also works for single domain networks. Users authenticate to the ASA using the same URLs mentioned above.

  When you use the User-Identity option, authentication is handled as follows:

  - If the user includes the domain in the login credentials, in the format DOMAIN\username, the ASA uses the domain to determine which AD server to use for authentication based on the domain mappings in the Identity Options policy. If no AAA server is mapped to the domain, the authentication attempt is rejected.
  - If the login credentials do not include an identifiable domain name (typically, if the \ character is not included in the username string), the ASA uses the AD server assigned to the default domain selected in the Identity Options policy. If no AAA server is mapped to the default domain, the authentication attempt will be rejected.

Tip Cut-through proxy works for IPv4 addresses only; IPv6 is not supported.

Related Topics

- Requirements for Identity-Aware Firewall Policies, page 13-3
- Configuring the Firewall to Provide Identity-Aware Services, page 13-7
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4
- Understanding How Users Authenticate, page 15-2

Step 1 Configure the Identity Options policy to specify all of the NetBIOS domains and their AD server groups, and the AD agent group, for the network, as described in Identifying Active Directory Servers and Agents, page 13-8.

Step 2 Do one of the following to open the AAA Rules Page, page 15-10:

- (Device view) Select **Firewall > AAA Rules** from the Policy selector.
- (Policy view) Select **Firewall > AAA Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 3 Create the following rules using the **Add Row** button. For detailed information about the fields in the Add AAA Rules dialog box, see Add and Edit AAA Rule Dialog Boxes, page 15-13.

Tip You can use more specific source, destination, or service specifications than the ones shown here.

- **Rule 1: Do not force users who have already authenticated to authenticate again.**
  - Select the **Authentication Action** and **User-Identity** options.
  - Action = Deny. For AAA authentication rules, "deny" means the user is not prompted for authentication, it does not mean the user's traffic is dropped.
  - Sources = any.
  - Users = all-auth-users.

    For users, **all-auth-users** means any user who has already authenticated to Active Directory, for which there is an IP mapping.
  - Destination = any.
  - Services = IP.
  - AAA Server Group = (no selection).
  - Interface = (your choice, typically inside interfaces).
- **Rule 2: Authenticate users who have not been authenticated yet.**
  - Select the **Authentication Action** and **User-Identity** options.
  - Action = Permit. This action requires matching users to authenticate.
  - User = all-unauth-users.

    In this case, **all-unauth-users** means any user who has not already authenticated to Active Directory.
  - All other options are identical to the first rule.

Collecting User Statistics

You can collect user statistics accounting information for identity-based firewall policies. These statistics are kept for users to which a firewall policy is applied based on username or user group membership.

Related Topics

- Requirements for Identity-Aware Firewall Policies, page 13-3
- Configuring the Firewall to Provide Identity-Aware Services, page 13-7
- Service Policy Rules Page, page 58-5
- Configuring Traffic Flow Objects, page 58-18

Step 1 Do one of the following:

- (Device view) Select an ASA device, then select **Platform > Service Policy Rules > IPS, QoS, and Connection Rules** from the Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Service Policy Rules > IPS, QoS, and Connection Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to add the rule, then click the **Add Row (+)** button below the table to start the Insert Service Policy Rule wizard.

Step 3 In step 1 of the wizard, select whether the rule will be Global or it will apply to specific interfaces or interface roles. Select Global if you want to collect statistics for users regardless of which interface their traffic passes through.

Click **Next**.

Step 4 In step 2, select the traffic class that defines the traffic for which you are collecting statistics. Select Use class-default if you want to collect statistics on all traffic. Otherwise, select Traffic Class and select the traffic flow object that defines the traffic matching attributes.

Click **Next**.

Step 5 In step 3, select the **User Statistics** tab.

- Select **Enable user statistics accounting**.
- Select the type of information you want to collect:
  - **Account for sent drop count**
  - **Account for sent packet, sent drop and received packet count**

Step 6 Click **Finish** to save your rule.

Filtering VPN Traffic with Identity-Based Rules

When you support remote access VPNs on an ASA, you configure user-sensitive access. You can also use identity-based rules to filter the traffic after validating the remote user access.

Before creating identity-based rules for VPN, understand the rules for VPN user names, to ensure that the rules use the correct domain name:

- If you use an Active Directory LDAP server group for authorization, and you configured that domain/server group in the Identity Options policy, the username is associated with the NetBIOS domain.
- For all other authorization mechanisms, the domain name for VPN users is LOCAL.

With this in mind, there are two methods you can use to filter the traffic on the VPN with identity-based ACL rules:

- Apply a VPN filter in the ASA Group Policy object. The filter applies to all users in the group. You can configure a VPN filter on the Connection Settings page in an ASA Group Policy object, which you use in a remote access connection policy. See ASA Group Policies Connection Settings, page 34-33.
- By default, VPN traffic bypasses interface access rules. You can change this behavior so that all VPN traffic must also pass through the interface access rules. If you take this approach, you must ensure that the interface rules are sensitive to your VPN traffic. To force VPN traffic to go through interface access rules, deselect the **Enable IPsec over Sysopt** option on the ISAKMP/IPsec tab of the RA VPN Global Settings policy. See Configuring VPN Global ISAKMP/IPsec Settings, page 26-33.

Monitoring Identity Firewall Policies

You can use Event Viewer to monitor identity-aware firewall policies the same way you would monitor other types of policies and events. The following are some tips to help you effectively monitor identity policies. For general information on using Event Viewer, see Chapter 69, "Viewing Events".

- There is a group of syslog messages that relate specifically to identity firewall: 746001-746019. You can find descriptions of these messages in the Syslog Message document for your ASA software version at http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html.

  The following messages are of particular concern:

  - **746004 and 746011**—These syslogs indicate that you have exceeded the supported number of references to user groups or users. You should consider changing your policies. For more information on these restrictions, see Requirements for Identity-Aware Firewall Policies, page 13-3.
  - **746003**—There was a failure in downloading user group or user mappings to IP address. The message explains the reason for the failure.
  - **746005**—The AD agent could not be reached. Ensure that the agent is functioning correctly and that there is a network path between the ASA and the agent.
  - **746010**—An update to the imported user or user group failed for the stated reason.
  - **746016**—DNS lookup for the fully-qualified domain name (FQDN) failed for the stated reason.
- Several existing syslog messages now include username or FQDN information. Event Viewer has two columns to display the information: Destination User Identity / FQDN and Source User Identity. Updated messages include:
  - 302005, 302006, 302013, 302014, 302016-302018, 302020, 302021.
  - 305005, 305006, 305009-305013.
  - 304001-304002 include identity information, but they are not parsed.
- You can filter on all identity-related syslog messages by creating a filter on Event Type and selecting the Identity Firewall Events folder.
- When you use the Go to Policy command on an event, as described in Looking Up a Security Manager Policy from Event Viewer, page 69-54, identity information is included in the lookup criteria. Note that identity information is not included in 106100, so policy lookup on that message cannot be sensitive to user identity.

Managing TrustSec Firewall Policies

Cisco TrustSec provides an access-control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec solution, enforcement devices utilize a combination of user attributes and end-point attributes to make role-based and identity-based access control decisions.

Cisco ASA, ISR, and ASR devices integrate with Cisco TrustSec to provide security group based policy enforcement. Access policies within the Cisco TrustSec domain are topology-independent, based on the roles of source and destination devices rather than on network IP addresses.

Note Cisco TrustSec Support for IOS feature is supported on the Cisco Integrated Services Router Generation 2 (ISR G2) only.

Security group awareness is integrated into several existing firewall rules; there is no unique TrustSec firewall policy. This chapter explains TrustSec firewall policies and how to implement them in the various policies that support security group awareness.

This chapter contains the following topics:

- Overview of TrustSec Firewall Policies, page 14-1
- Configuring TrustSec Firewall Policies, page 14-7
- Monitoring TrustSec Firewall Policies, page 14-17

Overview of TrustSec Firewall Policies

Traditionally, security features such as firewalls performed access control based on predefined IP addresses, subnets and protocols. However, with enterprises transitioning to borderless networks, both the technology used to connect people and organizations and the security requirements for protecting data and networks have evolved significantly. End points are becoming increasingly nomadic and users often utilize a variety of end points (for example, laptop versus desktop, smart phone, or tablet), which means that a combination of user attributes plus end-point attributes provide the key characteristics, in addition to existing 6-tuple based rules, that enforcement devices, such as switches and routers with firewall features or dedicated firewalls, can reliably use for making access control decisions.

As a result, the availability and propagation of end point attributes or client identity attributes have become increasingly important requirements to enable security solutions across the customers' networks, at the access, distribution, and core layers of the network and in the data center to name but a few examples.

Cisco TrustSec provides an access-control solution that builds upon an existing identity-aware infrastructure to ensure data confidentiality between network devices and integrate security access services on one platform. In the Cisco TrustSec solution, enforcement devices utilize a combination of user attributes and end-point attributes to make role-based and identity-based access control decisions.

Implementing Cisco TrustSec into your environment has the following advantages:

- Provides a growing mobile and complex workforce with appropriate and more secure access from any device
- Lowers security risks by providing comprehensive visibility of who and what is connecting to the wired or wireless network
- Offers exceptional control over activity of network users accessing physical or cloud-based IT resources
- Reduces total cost of ownership through centralized, highly secure access policy management and scalable enforcement mechanisms

For information about Cisco TrustSec, see http://www.cisco.com/go/trustsec.

This section contains the following topics:

- Understanding SGT and SXP Support in Cisco TrustSec, page 14-2
- Roles in the Cisco TrustSec Solution, page 14-3
- Security Group Policy Enforcement, page 14-3
- About Speaker and Listener Roles, page 14-6
- Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6

Understanding SGT and SXP Support in Cisco TrustSec

In the Cisco TrustSec solution, security group access transforms a topology-aware network into a role-based network, thus enabling end-to-end policies enforced on the basis of role-based access-control (RBAC). Device and user credentials acquired during authentication are used to classify packets by security groups. Every packet entering the Cisco TrustSec cloud is tagged with a security group tag (SGT). The tagging helps trusted intermediaries identify the source identity of the packet and enforce security policies along the data path.

An SGT can indicate a privilege level across the domain when the SGT is used to define a security group ACL. An SGT is assigned to a device through IEEE 802.1X authentication, web authentication, or MAC authentication bypass (MAB), which occurs with a RADIUS vendor-specific attribute. An SGT can be assigned statically to a particular IP address or to a switch interface. An SGT is passed along dynamically to a switch or access point after successful authentication.

The Security-group eXchange Protocol (SXP) is a protocol developed for Cisco TrustSec to propagate the IP-to-SGT mapping database across network devices that do not have SGT-capable hardware support to hardware that supports SGTs and security group ACLs. SXP, a control plane protocol, passes IP-SGT mappings from authentication points (such as legacy access layer switches) to upstream devices in the network.

The SXP connections are point-to-point and use TCP as the underlying transport protocol. SXP uses the well known TCP port number 64999 to initiate a connection. Additionally, an SXP connection is uniquely identified by the source and destination IP addresses.

Roles in the Cisco TrustSec Solution

To provide identity and policy-based access enforcement, the Cisco TrustSec solution includes the functionality:

- **Access Requestor (AR):** Access requestors are end-point devices that request access to protected resources in the network. They are primary subjects of the architecture and their access privilege depends on their Identity credentials.

  Access requestors include end-point devices such PCs, laptops, mobile phones, printers, cameras, and MACsec-capable IP phones.

- **Policy Decision Point (PDP):** A policy decision point is responsible for making access control decisions. The PDP provides features such as 802.1x, MAB, and Web authentication. The PDP supports authorization and enforcement through VLAN, DACL, and security group access (SGACL/SXP/SGT).

  In the Cisco TrustSec solution, the Cisco Identity Services Engine (ISE) acts as the PDP. The Cisco ISE provides identity and access control policy functionality.

- **Policy Information Point (PIP):** A policy information point is a source that provides external information (for example, reputation, location, and LDAP attributes) to policy decision points.

  Policy information points include devices such as Session Directory, Sensors IPS, and Communication Manager.

- **Policy Administration Point (PAP):** A policy administration point defines and inserts policies into authorization system. The PAP acts as an identity repository, by providing Cisco TrustSec tag to user identity mapping and Cisco Trustsec tag to server resource mapping.

  In the Cisco TrustSec solution, the Cisco Secure Access Control System (a policy server with integrated 802.1x and SGT support) acts as the PAP.

- **Policy Enforcement Point (PEP):** A policy enforcement point is the entity that carries out the decisions (policy rules and actions) made by the PDP for each AR. PEP devices learn identity information through the primary communication path that exists across networks. PEP devices learn the identity attributes of each AR from many sources, such as end-point agents, authorization servers, peer-enforcement devices, and network flows. In turn, PEP devices use SXP to propagate IP-SGT mappings to mutually-trusted peer devices across the network.

  Policy enforcement points include network devices such as Catalyst switches, routers, firewalls (specifically the ASA), servers, VPN devices, and SAN devices.

Security Group Policy Enforcement

Security policy enforcement is based on security group name. Compared to traditional IP-based policies configured on firewalls, identity-based policies are configured based on user and device identities. For example, mktg-contractor is allowed to access mktg-servers; mktg-corp-users are allowed to access mktg-server and corp-servers.

The benefits of this type of deployment include:

- User group and Resource is defined and enforced using a single object (SGT) - simplified policy management.
- User identity and resource identity are retained throughout the Cisco Trustsec capable switch infrastructure.

Figure 14-1 Security Group Name Based Policy Enforcement Deployment

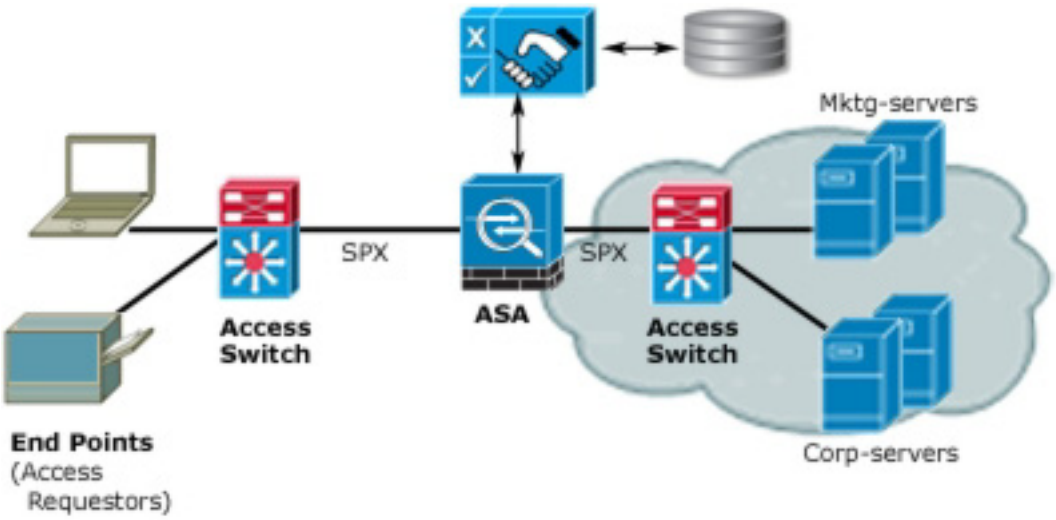

Implementing Cisco TrustSec allows you to configure security policies that support server segmentation and includes the following features:

- A pool of servers can be assigned an SGT for simplified policy management.
- The SGT information is retained within the infrastructure of Cisco Trustsec capable switches.
- The ASA, ISR, and ASR can leverage the IP-SGT mapping for policy enforcement across the Cisco TrustSec domain.
- Deployment simplification is possible because 802.1x authorization for servers is mandatory.

How the ASA Enforces Security Group Based Policies

Note User-based security policies and security-group based policies can coexist on the ASA. Any combination of network, user-based, and security-group based attributes can be configured in a security policy.

To configure the ASA to function with Cisco TrustSec, you must import a Protected Access Credential (PAC) file from the ISE.

Importing the PAC file to the ASA establishes a secure communication channel with the ISE. After the channel is established, the ASA initiates a PAC secure RADIUS transaction with the ISE and downloads Cisco TrustSec environment data (that is, the security group table). The security group table maps SGTs to security group names. Security group names are created on the ISE and provide user-friendly names for security groups.

Note For more information about the Cisco Identity Services Engine, see http://www.cisco.com/en/US/products/ps11640/index.html.

The first time that the ASA downloads the security group table, it walks through all entries in the table and resolves all the security group names included in security policies that have been configured on it; then the ASA activates those security policies locally. If the ASA cannot resolve a security group name, it generates a syslog message for the unknown security group name.

The following figure shows how a security policy is enforced in Cisco TrustSec.

Figure 14-2 Security Policy Enforcement

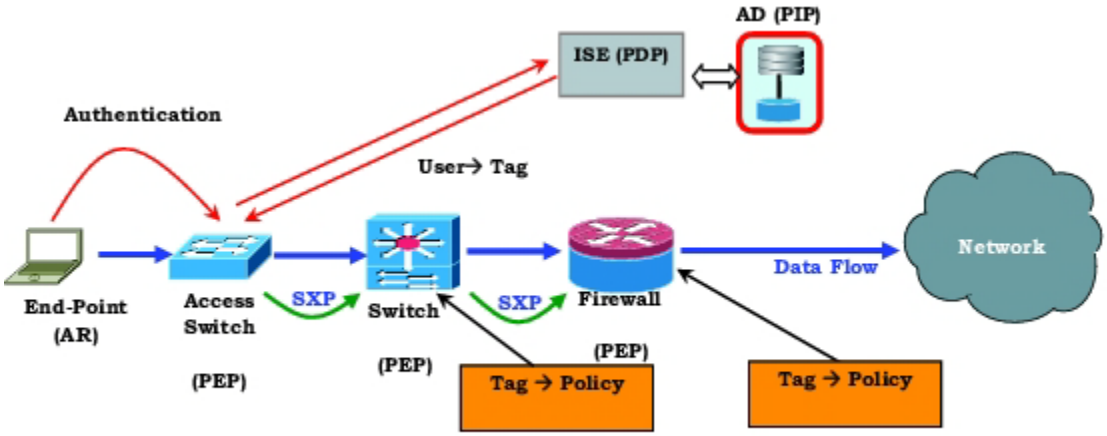

1. An end-point device connects to an access layer device directly or via remote access and authenticates with Cisco TrustSec.
2. The access layer device authenticates the end-point device with the ISE by using authentication methods such as 802.1X or web authentication. The end-point device passes role and group membership to classify the device into the appropriate security group.
3. The access layer device uses SXP to propagate the IP-SGT mapping to the upstream devices.
4. The ASA receives the packet and looks up the SGTs for the source and destination IP addresses using the IP-SGT mapping passed by SXP.

   If the mapping is new, the ASA records it in its local IP-SGT Manager database. The IP-SGT Manager database, which runs in the control plan, tracks IP-SGT mappings for each IPv4 or IPv6 address. The database records the source from which the mapping was learned. The peer IP address of the SXP connection is used as the source of the mapping. Multiple sources can exist for each IP-SGT mapping.

   If the ASA is configured as a Speaker, the ASA transmits all IP-SGT mappings to its SXP peers. See About Speaker and Listener Roles, page 14-6.
5. If a security policy is configured on the ASA with that SGT or security group name, the ASA enforces the policy. (You can create security policies on the ASA that include SGTs or security group names. To enforce policies based on security group names, the ASA needs the security group table to map security group names to SGTs.)

   If the ASA cannot find a security group name in the security group table and it is included in a security policy, the ASA considers the security group name unknown and generates a syslog message. After the ASA refreshes the security group table from the ISE and learns the security group name, the ASA generates a syslog message indicating that the security group name is known.

About Speaker and Listener Roles

The Security-group eXchange Protocol (SXP) is used to send and receive IP-SGT mappings to and from other network devices. Employing SXP allows security devices and firewalls to learn identity information from access switches without the need for hardware upgrades or changes. SXP can also be used to pass IP-SGT mappings from upstream devices (such as datacenter devices) back to the downstream devices.

When configuring an SXP connection to an SXP peer, you must designate the device as a Speaker or a Listener for that connection so that it can exchange identity information:

- Speaker mode—configures the device so that it can forward all active IP-SGT mappings to upstream devices for policy enforcement.
- Listener mode—configures the device so that it can receive IP-SGT mappings from downstream devices (SGT-capable switches) and use that information in creating policy definitions.

If one end of an SXP connection is configured as Speaker, then the other end must be configured as a Listener, and vice versa. If both devices on each end of an SXP connection are configured with the same role (either both as Speakers or both as Listeners), the SXP connection will fail and the device will generate a system log message.

Configuring the device to be both a Speaker and a Listener for an SXP connection can cause SXP looping, meaning that SXP data can be received by an SXP peer that originally transmitted it.

As part of configuring SXP, you configure an SXP reconcile timer. After an SXP peer terminates its SXP connection, the device starts a hold down timer. Only SXP peers designated as Listener devices can terminate a connection. If an SXP peer connects while the hold down timer is running, the device starts the reconcile timer; then, the device updates the IP-SGT mapping database to learn the latest mappings.

Prerequisites for Integrating an ASA with Cisco TrustSec

Before configuring the ASA to integrate with Cisco TrustSec, you must perform the following prerequisites:

- Register the ASA with the ISE.
- Create a security group for the ASA on the ISE.
- Generate the PAC file on the ISE to import into the ASA.

Registering the ASA with the ISE

The ASA must be configured as a recognized Cisco TrustSec network device in the ISE before the ASA can successfully import a PAC file.

1. Log into the ISE.
2. Choose **Administration > Network Devices > Network Devices**.
3. Click **Add**.
4. Enter the IP address of the ASA.
5. When the ISE is being used for user authentication in the Cisco TrustSec solution, enter a shared secret in the Authentication Settings area.

   When you configure the AAA sever on the ASA, provide the shared secret you create here on the ISE. The AAA server on the ASA uses this shared secret to communicate with the ISE.

6. Specify a device name, device ID, password, and a download interval for the ASA. See the ISE documentation for the details to perform these tasks.

Creating a Security Group on the ISE

When configuring the ASA to communicate with the ISE, you specify a AAA server. When configuring the AAA server on the ASA, you must specify a server group.

The security group must be configured to use the RADIUS protocol.

1. Log into the ISE.
2. Choose **Policy > Policy Elements > Results > Security Group Access > Security Group**.
3. Add a security group for the ASA. (Security groups are global and not ASA specific.)

   The ISE creates an entry under Security Groups with a tag.
4. Under the Security Group Access section, configure a device ID credentials and password for the ASA.

Generating the PAC

Before generating the PAC file, you must have registered the ASA with the ISE.

1. Log into the ISE.
2. Choose **Administration > Network Resources > Network Devices**.
3. From the list of devices, select the ASA device.
4. Under the Security Group Access (SGA), click **Generate PAC**.
5. To encrypt the PAC file, enter a password.

   The password (or encryption key) you enter to encrypt the PAC file is independent of the password that was configured on the ISE as part of the device credentials.

The ISE generates the PAC file. The ASA can import the PAC from flash or from a remote server via TFTP, FTP, HTTP, HTTPS, or SMB. (The PAC does not have to reside on the ASA flash before you can import it.)

Configuring TrustSec Firewall Policies

Security group awareness is integrated into several existing firewall rules; there is no unique TrustSec firewall policy. Additionally, supporting tools have been updated to work on TrustSec firewall policies. For example, you can search for rules that include a specific Security Group using the Find and Replace tool.

The topics in this section explain the various procedures for integrating security group awareness into firewall policies.

This section contains the following topics:

- Configuring Cisco TrustSec Services, page 14-8
- Creating Security Group Objects, page 14-14
- Selecting Security Groups in Policies, page 14-16
- Configuring TrustSec-Based Firewall Rules, page 14-17

Configuring Cisco TrustSec Services

This procedure explains how to enable and configure Cisco TrustSec in Cisco Security Manager and on the required security devices.

Before You Begin

Before configuring an ASA to integrate with Cisco TrustSec, you must meet the prerequisites explained in Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6.

To configure Cisco TrustSec, perform the following tasks:

Step 1 Configure communication between Cisco Security Manager and the Cisco Identity Services Engine (ISE). See ISE Settings Page, page 11-56.

Note Security Manager supports communications with only one ISE appliance/server for fetching and resolving security group names and tags.

Step 2 Enable and set the default values for the Security Exchange Protocol (SXP). See Configuring Security Exchange Protocol (SXP) Settings, page 14-8.

Step 3 Add SXP connection peers for the Cisco TrustSec architecture. See Defining SXP Connection Peers, page 14-12.

Step 4 (ASA 9.3.1+ devices only) Configure Security Group Tagging options. See Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41.

Step 5 (ASA 9.3.1+ devices only) Configure Security Group Tagging for VPN sessions. See ASA Group Policies SSL VPN Full Client Settings, page 34-19.

Step 6 Configure the Security Policy. See Configuring TrustSec-Based Firewall Rules, page 14-17.

Step 7 Monitor the TrustSec firewall system. See Monitoring TrustSec Firewall Policies, page 14-17.

Configuring Security Exchange Protocol (SXP) Settings

Use the SXP Settings page to enable the Security Exchange Protocol (SXP) on your security device and to configure SXP settings for the device.

Note All settings are available on the SXP Settings page whether you access that page from Policy view or from Device view for a specific device type. If you configure a setting that is not supported on a particular device, you will receive a validation warning and the unsupported CLI will not be generated for the device.

Navigation Path

- (Device view) Select the security device, then select **TrustSec > SXP Settings** from the Policy selector.
- (Policy view) Select **TrustSec > SXP Settings** from the Policy selector. Select an existing policy or create a new one.

Related Topics

- Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6
- Defining SXP Connection Peers, page 14-12

Field Reference

Table 14-1 SXP Settings Page

Element	Description
Enable SGT Exchange Protocol (SXP)	Whether to enable the Security Exchange Protocol on the device. The default is disabled
Retry Timer	The default time interval between attempts to set up new SXP connections between SXP peers. Enter the retry timer value as a number of seconds in the range of 0 to 64000 seconds. If you specify 0 seconds, the timer never expires and the device will not attempt to connect to SXP peers. By default, the timer value is 120 seconds. The device will continue to attempt to connect to new SXP peers until a successful connection is made. The retry timer is triggered as long as there is one SXP connection on the device that is not up. When the retry timer expires, the device goes through the connection database and if the database contains any connections that are off or in a "pending on" state, the device restarts the retry timer.
Reconcile Timer	The reconcile timer value as a number of seconds in the range of 0 to 64000 seconds. By default, the timer value is 120 seconds. After an SXP peer terminates its SXP connection, the security device starts a hold down timer. If an SXP peer reconnects while the hold down timer is running, the device starts the reconcile timer; then, the device updates the SXP mapping database to learn the latest mappings. When the reconcile timer expires, the device scans the SXP mapping database to identify stale mapping entries (entries that were learned in a previous connection session). The device marks these connections as obsolete. When the reconcile timer expires, the device removes the obsolete entries from the SXP mapping database. **Note** Setting the reconciliation period to 0 seconds disables the timer and causes all entries from the previous connection to be removed.
Network Map	The Network Map argument specifies the maximum number of subnet IP hosts from 0 to 65,535 that can be bound to SGTs and exported to the SXP listener. The default is 0 (no expansions performed).

Table 14-1 SXP Settings Page (continued)

Element	Description
Server Group Name (not applicable for IOS-XE)	Enter or select the name of the security group created on the ISE for the device. **Note** If you choose to select a server group, you are also give the option to add a AAA Server group. The server group name you specify here must match the name of the security group created on the ISE for the device. If these two group names do not match, the device will not be able to communicate with the ISE. Contact your ISE administrator if you do not have this information.
CTS Server Settings (IOS/IOS-XE Only)	
Log Binding Changes	Whether to enable logging for IP-to-SGT binding changes causing SXP syslogs (sev 5 syslog) to be generated whenever a change to IP-to-SGT binding occurs (add, delete, change). These changes are learned and propagated on the SXP connection. This logging function is disabled by default.
Enable Cache Cache NV Storage (not applicable for IOS-XE)	Whether to enable caching of TrustSec authorization and environment data information to DRAM and NVRAM. To have DRAM cache updates written to non-volatile storage and to enable the DRAM cache to be initially populated from non-volatile storage when the device boots, select the desired file system from the Cache NV Storage list. Options include: • flash • flash0 • flash1 • flash2 • disk0 • disk1 • disk2
CTS SGT Number	Enter a number from 1-65533 to manually assign a Security Group Tag (SGT) number for this device.
Server Dead Time (not applicable for IOS-XE)	Specifies how long a server in the group should not be selected for service once it has been marked as dead. The default is 20 seconds; the range is 1 to 864000.

Table 14-1 SXP Settings Page (continued)

Element	Description
Load Balance (not applicable for IOS-XE)	Whether to configure RADIUS server group load balancing. When Load Balance is enabled, the following options can be specified: Batch size—The number of transactions to be assigned per batch. The default transactions is 25. **Note** Changes in batch size may impact CPU load and network throughput. As batch size increases, CPU load decreases and network throughput increases. However, if a large batch size is used, all available server resources may not be fully utilized. As batch size decreases, CPU load increases, and network throughput decreases. It is recommended that the default batch size, 25, be used because it is optimal for high throughput, without adversely impacting CPU load. Ignore Preferred-Server—Instructs the device not to try to use the same server throughout a session.
SGT Rolebased Map table (ASA 9.3(1)+, IOS15.2(2)T+, and IOS-XE3.5.x(15.2(1)S)+ only)	Use the SGT Rolebased Map table to manually map Security Group Tag (SGT) numbers to individual IP addresses or host objects. You can do the following: • To add an entry, click the **Add Row (+)** button and fill in the Add Connection Peer dialog box. See Add/Edit SGT Role Dialog Box, page 14-11. • To edit an entry, select it and click the **Edit Row (pencil)** button. • To delete an entry, select it and click the **Delete Row (trash can)** button.

Add/Edit SGT Role Dialog Box

Use the Add/Edit SGT Role dialog box to manually map Security Group Tag (SGT) numbers to individual IP addresses or host objects.

Navigation Path

- (Device view) Select an ASA, ISR, or ASR, then select **TrustSec > SXP Settings** from the Policy selector.
  - To add an entry, click the **Add Row (+)** button beneath the SGT Rolebased Map table.
  - To edit an entry, select it and click the **Edit Row (pencil)** button beneath the SGT Rolebased Map table.
- (Policy view) Select **TrustSec > SXP Settings** from the Policy selector. Select an existing policy or create a new one.
  - To add an entry, click the **Add Row (+)** button beneath the SGT Rolebased Map table.
  - To edit an entry, select it and click the **Edit Row (pencil)** button beneath the SGT Rolebased Map table.

Related Topics

- About Speaker and Listener Roles, page 14-6

- Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6

Field Reference

Table 14-2 Add/Edit SGT Role dialog box

Element	Description
IP Address	The IPv4 address of the host for which you want to manually assign a Security Group Tag (SGT) number. You can enter an IP address or the name of a host object, or click **Select** to select the object from a list or to create a new one.
CTS SGT Number	The Security Group Tag (SGT) number to assign to the specified host/IP address. Valid security tag numbers are 2-65519 for ASA 9.3(1)+ and 0-65533 for ISR and ASR.

Defining SXP Connection Peers

The Security-group eXchange Protocol (SXP) is a protocol developed for Cisco TrustSec to propagate the IP-to-SGT mapping database across network devices that do not have SGT-capable hardware support to hardware that supports SGTs and security group ACLs. SXP, a control plane protocol, passes IP-SGT mappings from authentication points (such as legacy access layer switches) to upstream devices in the network. SXP connections between peers are point-to-point and use TCP as the underlying transport protocol.

Related Topics

- Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6
- About Speaker and Listener Roles, page 14-6
- Configuring Security Exchange Protocol (SXP) Settings, page 14-8

Step 1 Do one of the following:

- (Device view) Select an ASA, ISR, or ASR, then select **TrustSec > SXP Connection Peers** from the Policy selector.
- (Policy view) Select **TrustSec > SXP Connection Peers** from the Policy selector. Select an existing policy or create a new one.

Step 2 In **Default Source**, enter the default local IP address for SXP connections. You can enter an IP address or the name of a network/host object, or click Select to select the object from a list or to create a new one. The IP address can be an IPv4 or IPv6 address.

Note The device determines the local IP address for an SXP connection as the outgoing interface IP address that is reachable by the peer IP address. If the configured local address is different from the outgoing interface IP address, the device cannot connect to the SXP peer and generates a system log message.

Step 3 In **Default Password** and **Confirm**, enter the default password for TCP MD5 authentication with SXP peers. By default, SXP connections do not have a password set.

You can specify the password as an encrypted string up to 162 characters or an ASCII key string up to 80 characters.

Step 4 Configure the SXP Peers:

You can do the following:

- To add an entry, click the **Add Row (+)** button and fill in the Add Connection Peer dialog box. See Add/Edit Connection Peer Dialog Box, page 14-13.
- To edit an entry, select it and click the **Edit Row (pencil)** button.
- To delete an entry, select it and click the **Delete Row (trash can)** button.

Step 5 Click **Save** to save your changes.

Add/Edit Connection Peer Dialog Box

Use the Add/Edit Connection Peer dialog box to define the settings for an SXP Connection.

Note All settings are available on the Add/Edit Connection Peer dialog box whether you access that dialog box from Policy view or from Device view for a specific device type. If you configure a setting that is not supported on a particular device, you will receive a validation warning and the unsupported CLI will not be generated for the device.

Navigation Path

- (Device view) Select an ASA, ISR, or ASR, then select **TrustSec > SXP Connection Peers** from the Policy selector.
  - To add an entry, click the **Add Row (+)** button.
  - To edit an entry, select it and click the **Edit Row (pencil)** button.
- (Policy view) Select **TrustSec > SXP Connection Peers** from the Policy selector. Select an existing policy or create a new one.
  - To add an entry, click the **Add Row (+)** button.
  - To edit an entry, select it and click the **Edit Row (pencil)** button.

Related Topics

- About Speaker and Listener Roles, page 14-6
- Prerequisites for Integrating an ASA with Cisco TrustSec, page 14-6

Field Reference

Table 14-3 Add Connection Peer dialog box

Element	Description
Peer IP Address	The IPv4 or IPv6 address of the SXP peer. The peer IP address must be reachable from the outgoing interface. You can enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one.

Table 14-3 Add Connection Peer dialog box (continued)

Element	Description
Source IP Address	(Optional) The local IPv4 or IPv6 address of the SXP connection. Specifying the source IP address is optional, however, specifying it safeguards misconfiguration. You can enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one. **Note** You cannot configure the Source IP Address and Peer IP Address with the same address. Also, you cannot use an IPv4 address with one field and an IPv6 address with the other.
Password	Whether to use the authentication key for the SXP connection. Select from the following values: • default—Use the default password configured for SXP connections. See Defining SXP Connection Peers, page 14-12. • none—Do not use a password for the SXP connection.
Mode	The mode of the SXP connection. Select from the following values: • local—Use the local SXP device. • peer—Use the peer SXP device.
Role	Whether the device functions as a Speaker or Listener for the SXP connection: • listener—The device can receive IP-SGT mappings from downstream devices. • speaker—The device can forward IP-SGT mappings to upstream devices. See About Speaker and Listener Roles, page 14-6.
Hold Time Min Only applicable on IOS and IOS-XE	The minimum length of the hold-time period in seconds for the speaker or listener device.
Hold Time Max Only applicable on IOS and IOS-XE	The maximum length of the hold-time period in seconds for the speaker or listener device. A hold-time maximum-period value is required only when you use the following option combinations: peer speaker and local listener. In other instances, only a hold-time minimum-period value is required.

Creating Security Group Objects

You can create security group object groups for use in features that support Cisco TrustSec by including the group in an extended ACL, which in turn can be used in an access rule, for example.

When integrated with Cisco TrustSec, the security device downloads security group information from the Cisco Identity Services Engine (ISE). The ISE acts as an identity repository, by providing Cisco TrustSec tag to user identity mapping and Cisco TrustSec tag to server resource mapping. You provision and manage security group access lists centrally on the ISE.

However, the device might have localized network resources that are not defined globally that require local security groups with localized security policies. Local security groups can contain nested security groups that are downloaded from the ISE. The security device consolidates local and central security groups.

To create local security groups on the device, you create a local security object group. A local security object group can contain one or more nested security object groups or Security IDs or security group names. Users can also create a new Security ID or security group name that does not exist on the device.

You can use the security object groups you create to control access to network resources. You can use the security object group as part of an access group or service policy.

Tips

- Use of these objects is supported on ASA 9.0(1)+, IOS 15.2(2)T+, and IOS-XE 3.5.x(15.2(1)S)+ only.
- Although IOS does not support security group object groups, Security Manager allows the use of security group object groups for ISR devices and will automatically expand the object groups to their tag numbers during deployment.
- You must configure the TrustSec policy on the device to enable the use of these objects.
- You can create security group objects when defining policies or objects that use this object type. For more information, see Selecting Security Groups in Policies, page 14-16.

Related Topics

- Selecting Security Groups in Policies, page 14-16
- Creating Policy Objects, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **Security Group** from the Object Type selector.

Step 3 Right-click in the work area, then select **New Object** to open the Add Security Group dialog box.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Add and remove items in the **Members in Group** list to identify the users and user groups defined in the object.

To populate the list, do any combination of the following:

- In **Available Security Group**, select an existing object and click the **Add >>** button between the lists.
- In **Search name/tag**, select a security group from the ISE server configured in the ISE Settings administrative options. You must configure the settings before you can select a name or tag (see ISE Settings Page, page 11-56).

  To find a security group, enter a search string. Then, click **Search** to find matches. A name is considered a match if the string appears anywhere within the security group name.

  To add the security group, select it in the list and click the **Add >>** button between the lists.
- In **Type in comma separated (Name or Tag)**, first select the type of entry you are making, Name or Tag. Type in a valid security group name or tag number, then click the **Add >>** button between the lists. Separate multiple names or tags with commas; they are added as separate lines in the members list. When adding multiple names or tags, avoid adding spaces before or after the comma.

Valid security tag numbers are 0-65533 for ASA 9.3+, 1-65533 for ASA versions less than 9.3, 0-65533 for ISR, and 1-65535 for ASR.

- To remove an item from the object, select it in the Members list and click the **<< Remove** button between the lists.

Note Although IOS and IOS-XE only support security group tags, Security Manager allows the use of security group names for ISR and ASR devices, either directly or as part of a security group object group, and will automatically convert security group names to their tag number during deployment. However, if the ISE is not reachable, Security Manager will not be able to resolve the security group names. In such situations, Security Manager will not generate the rule for the configuration and an appropriate warning/error message will be displayed.

Step 6 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 7 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 8 Click **OK** to save the object.

Selecting Security Groups in Policies

In any policy or policy object that allows the specification of security groups, whether directly or through the selection of a TrustSec security group object, you can click the Select button next to the Security Groups field to help you enter the information.

In the Security Group Selector dialog box, you can define the content of the Security Groups field by populating the **Members in Group** list. To populate the list, do any combination of the following:

- In **Available Security Group**, select an existing object and click the **Add >>** button between the lists. If the desired object does not exist, you can click the **Add (+)** button below the list to create a new object. You can also select an object and click the **Edit (pencil)** button to modify it or to examine its contents.
- In **Search name/tag**, select a security group from the ISE server configured in the ISE Settings administrative options. You must configure the settings before you can select a name or tag (see ISE Settings Page, page 11-56).

  To find a security group, enter a search string. Then, click **Search** to find matches. A name is considered a match if the string appears anywhere within the security group name.

  To add the security group, select it in the list and click the **Add >>** button between the lists.
- In **Type in comma separated (Name or Tag)**, first select the type of entry you are making, Name or Tag. Type in a valid security group name or tag number, then click the **Add >>** button between the lists. Separate multiple names or tags with commas; they are added as separate lines in the members list. When adding multiple names or tags, avoid adding spaces before or after the comma.

  Valid security tag numbers are 0-65533 for ASA 9.3+, 1-65533 for ASA versions less than 9.3, 0-65533 for ISR, and 1-65535 for ASR.
- To remove an item from the object, select it in the Members list and click the **<< Remove** button between the lists.

Configuring TrustSec-Based Firewall Rules

Security group awareness is integrated into the access control entries, or rules, in the ACLs used to provide firewall services. Because the feature is integrated into the ACL, the techniques for adding security group awareness to a firewall policy are the same for all types of firewall policy. This topic provides general guidance on how to incorporate security group awareness into your existing policies, and directs you to more specific information on configuring each type of policy that supports security groups.

Firewall Policies That Support Security Groups

For ASA 9.0.1+ only, you can configure security groups for the following policy types:

- AAA Rules—Select **Firewall > AAA Rules** and see Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4.
- Access Rules—Select **Firewall > Access Rules** and see Configuring Access Rules, page 16-7.
- Inspection Rules—Select **Firewall > Inspection Rules** and see Configuring Inspection Rules, page 17-5.
- Policies that use extended ACL policy objects—Several firewall policies use extended ACL policy objects to define traffic matching criteria instead of incorporating a rule table directly in the policy. You can configure extended ACL policy objects to include security group specifications (see Creating Extended Access Control List Objects, page 6-54). You can then use these extended ACL objects in the following policies:
  - Botnet Traffic Filter Rules—Select **Firewall > Botnet Traffic Filter Rules** and see Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6. You can use security groups as part of the traffic classification for Enable and Drop rules.
  - IPS, QoS, and Connection Rules (service policy rules)—Select **Platform > Service Policy Rules > IPS, QoS, and Connection Rules** and see Service Policy Rules Page, page 58-5.

    Traffic match criteria in this policy is based on extended ACL policy objects that are incorporated into traffic flow policy objects. You must select one of the options for specifying an ACL in the traffic flow object to incorporate security group traffic classification. For more information, see Configuring Traffic Flow Objects, page 58-18.

For devices running IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+, you can configure security groups for Zone-based Firewall Rules (**Firewall > Zone Based Firewall Rules**). For more information, see Adding Zone-Based Firewall Rules, page 21-13.

Monitoring TrustSec Firewall Policies

You can use Event Viewer to monitor TrustSec firewall policies the same way you would monitor other types of policies and events. The following are some tips to help you effectively monitor identity policies. For general information on using Event Viewer, see Chapter 69, “Viewing Events”.

- There are groups of syslog messages that relate specifically to Cisco TrustSec: 766001-766020, 766201-766205, 766251-766254, and 766301-766313. You can find descriptions of these messages in the Syslog Message document for your ASA software version at http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html.
- Event Viewer has the following columns to display TrustSec information: TrustSec Security Group Name, TrustSec Security Group Tag, SXP Connection Source IP, SXP Connection Failure Reason, SXP Peer IP, SXP Peer Connection Failure Reason.

- You can filter on all identity-related syslog messages by creating a filter on Event Type and selecting the All Firewall Events > Trustsec Events folder.

Managing Firewall AAA Rules

You can use Authentication, Authorization, and Accounting (AAA) rules to control access to network resources based on user privileges rather than by IP addresses. If you configure authentication rules, users must enter a username and password whenever they attempt to access a network behind the protected device. Once authenticated, you can further require that the user account be checked to ensure the user is authorized for network access. Finally, you can use accounting rules to track access for billing, security, or resource allocation purposes.

AAA rule configuration is complex and requires that you configure more than just the AAA rules policy. The following topics explain AAA rules in greater detail and include procedures that explain not only the AAA rules policy configuration but also what you must configure in related policies:

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4
- Configuring AAA Rules for IOS Devices, page 15-7
- AAA Rules Page, page 15-10
- AAA Firewall Settings Policies, page 15-20

The following topics can help you with general rule table usage:

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19

Understanding AAA Rules

You can use Authentication, Authorization, and Accounting (AAA) rules to control access to network resources based on user privileges rather than by IP addresses. AAA rules provide a different type of control compared to traditional access rules; where access rules allow you to control which IP addresses and services are allowed, AAA rules allow you to configure ACLs for each user to define the authorization available on a user basis, regardless of the IP address from which the user connects. (These per-user ACLs are configured in the AAA server, not in the AAA rule defined on the device.)

AAA rules policies differ from other device platform AAA policies in that AAA rules apply to traffic that is passing through the device, not to traffic directed specifically at the device. By using AAA rules, you can control entry into, or out of, a network. This might be useful if you have a network segment that

is high security, where you want to carefully control access. AAA rules are also useful for circumstances where you need to maintain per-user accounting records for billing, security, or resource allocation purposes.

The AAA rules policy actually configures three separate types of rule, and the configuration of these rules differs significantly between IOS devices on the one hand and ASA, PIX, and FWSM devices on the other hand. For IOS devices, these policies define what is called authentication proxy admission control. When creating shared AAA rules, create separate rules for these types of devices. Following are the types of rules you can configure with AAA rules:

- Authentication rules—Authentication rules control basic user access. If you configure an authentication rule, users must log in if their connection request goes through the device on which the rule is defined. You can force users to log in for HTTP, HTTPS, FTP, or Telnet connections. For ASA, PIX, and FWSM devices, you can control other types of services, but users must first authenticate using one of the supported protocols before other types of traffic are allowed.

  The device recognizes these traffic types only on the default ports: FTP (21), Telnet (23), HTTP (80), HTTPS (443). If you map these types of traffic to other ports, the user will not be prompted, and access will fail.

- Authorization rules—You can define an additional level of control over and above authentication. Authentication simply requires that users identify themselves. After authentication is successful, an authorization rule can query the AAA server to determine if the user has sufficient privileges to complete the attempted connection. If authorization fails, the connection is dropped.
  - For ASA, PIX, and FWSM devices, you define authorization rules directly in the AAA rules policy; if you require authorization for traffic that does not also require authentication, the unauthenticated traffic is always dropped. If you use RADIUS servers for authentication, authorization is automatically performed and authorization rules are not necessary. If you configure authorization rules, you must use a TACACS+ server.
  - For IOS devices, to configure authorization, you must configure an authorization server group in the **Firewall > Settings > AAA** policy; authorization is done for any traffic that is subject to authentication. You can use TACACS+ or RADIUS servers.

- Accounting—You can define accounting rules even if you do not configure authentication or authorization. If you do configure authentication, accounting records are created for each user, so that you can identify the specific user who made the connection. Without user authentication, accounting records are based on IP address. You can use TACACS+ or RADIUS servers for accounting.
  - For ASA, PIX, and FWSM devices, you define accounting rules directly in the AAA rules policy. You can perform accounting for any TCP or UDP protocol.
  - For IOS devices, to configure accounting, you must configure an accounting server group in the **Firewall > Settings > AAA** policy; accounting is done for any traffic that is subject to authentication.

Understanding How Users Authenticate

When you create AAA rules to require that users authenticate when trying to make connections through a device, users will be prompted to supply credentials: a username and password. These credentials must be defined in a AAA server or in the local database configured on the device.

Users are prompted only for HTTP, HTTPS, FTP, and Telnet connections (if you configure those protocols to require authentication). For ASA, PIX, and FWSM devices, you can also require authentication for other protocols; however, users are not prompted for them, and so they must first attempt one of the four supported protocols and successfully authenticate before completing connections of any other protocol that requires authentication.

Tip For ASA, PIX, and FWSM devices, if you do not want to allow HTTP, HTTPS, Telnet, or FTP through the security appliance but want to authenticate other types of traffic, you can require that the user authenticate with the security appliance directly using HTTP or HTTPS by configuring the interface to use interactive authentication (using the **Firewall > Settings > AAA Firewall** policy). The user would then authenticate with the appliance before trying other connections, using one of the following URLs, where *interface_ip* is the IP address of the interface and *port* is optionally the port number, if you specify a non-default port for the protocol in the interactive authentication table:
http://*interface_ip*[**:***port*]**/netaccess/connstatus.html** or
https://*interface_ip*[**:***port*]**/netaccess/connstatus.html**.

When attempting a connection through the device, the user is prompted based on the protocol:

- HTTP—The device prompts the user with a web page to provide username and password. The user is prompted repeatedly until successfully authorized. After the user authenticates correctly, the device redirects the user to the original destination. If the destination server also has its own authentication, the user enters another username and password.

  For ASA, PIX, and FWSM devices, the security appliance uses basic HTTP authentication by default, and provides an authentication prompt. You can improve the user experience by configuring the interface for interactive authentication and specifying redirect for HTTP traffic. This redirects the user to a web page hosted on the appliance for authentication. To configure an interface to use interactive authentication, add the interface to the Interactive Authentication table on the **Firewall > Settings > AAA Firewall** policy (see AAA Firewall Settings Page, Advanced Setting Tab, page 15-20). Ensure that you select the HTTP and Redirect options when adding the interface.

  You might want to continue to use basic HTTP authentication if: you do not want the security appliance to open listening ports; if you use NAT on a router and you do not want to create a translation rule for the web page served by the security appliance; basic HTTP authentication might work better with your network. For example non-browser applications, like when a URL is embedded in email, might be more compatible with basic authentication.

  However, when using basic HTTP authentication, if the user is going to an HTTP server that requires authentication, the same username and password used to authenticate with the appliance is sent to the HTTP server. Thus, login to the HTTP server fails unless the same username and password are used by the ASA and HTTP server. To avoid this problem, you must configure a virtual HTTP server on the ASA. You can configure a virtual HTTP server using the **Firewall > Settings > AAA Firewall** policy (see AAA Firewall Settings Page, Advanced Setting Tab, page 15-20).

Tip In HTTP authentication, the username and password are transmitted in clear text. You can prevent this by selecting the **Use Secure HTTP Authentication** option on the **Firewall > Settings > AAA Firewall** policy. This option ensures that credentials are encrypted.

- HTTPS—The user experience for HTTPS is the same as for HTTP; the user is prompted until successfully authorized, and then redirected to the original destination.

For ASA, PIX, and FWSM devices, the security appliance uses a custom login screen. Like with HTTP, you can configure the interface to use interactive authentication, in which case HTTPS connections use the same authentication page as HTTP connections. You must configure the interface separately for HTTPS redirection; use the **Firewall > Settings > AAA Firewall** policy.

For IOS devices, HTTPS connections are authenticated only if you enable SSL on the device and your AAA rules require HTTP authentication proxy. This configuration is explained in Configuring AAA Rules for IOS Devices, page 15-7.

- FTP—The device prompts once for authentication. If authentication fails, the user must retry the connection.

  When prompted, the user can enter the username required for device authentication followed by an at sign (@) and then the FTP username (name1@name2). For the password, the user would then enter the device authentication password followed by an at sign (@) and then the FTP password (password1@password2). For example, enter the following text.

  name> asa1@partreq
  password> letmein@he110

  For IOS devices, this method of entering both the device and FTP credentials is required. For ASA, PIX, and FWSM devices, this feature is useful when you have cascaded firewalls that require multiple logins. You can separate several names and passwords by multiple at signs (@).

- Telnet—The device prompts several times for authentication. After a number of failed attempts, the user must retry the connection. After authentication, the Telnet server prompts for its username/password. You can configure a virtual Telnet server using the **Firewall > Settings > AAA Firewall** policy (see AAA Firewall Settings Page, Advanced Setting Tab, page 15-20).

Configuring AAA Rules for ASA, PIX, and FWSM Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

When you configure AAA rules for an ASA, PIX, or FWSM device, you are configuring policies that define who is allowed to make HTTP, HTTPS, FTP, and Telnet connections through (not to) the device. To fully configure network access authentication, you need to configure several policies, not just the AAA rules policy.

The following procedure covers all policies you would need to configure to supply full authentication, authorization, and accounting support for network access authentication. You do not need to configure options for features you do not need.

Related Topics

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Creating a New Shared Policy, page 5-54
- Modifying Policy Assignments in Policy View, page 5-54
- Understanding Networks/Hosts Objects, page 6-80
- Understanding Interface Role Objects, page 6-73
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

- Understanding AAA Server and Server Group Objects, page 6-27
- Understanding Interface Role Objects, page 6-73

Step 1 Do one of the following to open the AAA Rules Page, page 15-10:

- (Device view) Select **Firewall > AAA Rules** from the Policy selector.
- (Policy view) Select **Firewall > AAA Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add and Edit AAA Rule Dialog Boxes, page 15-13.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10.

Step 3 Configure the rule. Following are the highlights of what you typically need to decide. For specific information on configuring the fields, see Add and Edit AAA Rule Dialog Boxes, page 15-13.

- Authentication (with or without User-Identity), Authorization, or Accounting Action—Select the options applicable for this rule. Authentication prompts the user for a username and password when attempting HTTP, HTTPS, FTP, or Telnet access. Authorization is an additional level, where after the user authenticates, the AAA server is checked to ensure that the user is authorized for that type of access. Accounting generates usage records in the AAA server and can be used for billing, security, or resource allocation purposes. You can generate accounting information for any TCP or UDP traffic.

  When you select Authentication, you can also select User-Identity (ASA 8.4(2+) only). This option indicates that the ASA should use the Active Directory servers configured in the identity-firewall domain mappings to authenticate users (see Identifying Active Directory Servers and Agents, page 13-8). If the user enters a domain name, the AD server associated with the domain is queried. Otherwise, the AD server associated with the default domain is queried. When you select User-Identity, and you do not select Authorization or Accounting, do not specify a AAA server group.
- Permit or Deny—Whether you are subjecting the identified traffic to AAA control (permit) or you are exempting it from AAA control (deny). Any denied traffic is not prompted for authentication and is allowed to pass unauthenticated, although your access rules might drop the traffic.
- Source and Destination addresses—If the rule should apply no matter which addresses generated the traffic or their destinations, use “All-Addresses” as the source or destination. If the rule is specific to a host or network, enter the addresses or network/host objects. For information on the accepted address formats, see Specifying IP Addresses During Policy Definition, page 6-87.
- Source and Destination Security Groups (ASA 9.0+ only)—You can specify TrustSec security groups used to filter traffic in addition to the source and destination addresses. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups.
- Source Users (ASA 8.4.2+ only)—You can further define the traffic source by specifying Active Directory (AD) user names (in the format NetBIOS_DOMAIN\username), user groups (NetBIOS_DOMAIN\\user_group), or identity user group objects that define the names and groups. The user specification is conjoined to the source address to limit the match to user addresses within the source address range. For more information, see Configuring Identity-Based Firewall Rules, page 13-21 and Creating Identity User Group Objects, page 13-19.

- Services—You can specify any type of service for authentication and authorization rules; however, the user is prompted to authenticate only for HTTP, HTTPS, FTP, and Telnet connections. Thus, if you specify something other than these services, the user must first attempt one of these connections and successfully authenticate (and be authorized, if you include that action) before any other types of connections are allowed. For accounting rules, you can specify any TCP or UDP service (or simply TCP and UDP themselves), if you want to account for all types of traffic.
- AAA Server Group—The AAA server group policy object to be used for authentication, authorization, or accounting. If the rule applies more than one of these actions, the server group must support all selected actions. For example, only TACACS+ servers can provide services for authorization rules (although using RADIUS for authentication rules automatically includes RADIUS authorization), and only TACACS+ and RADIUS servers can provide accounting services. If you want to use different server groups for particular actions, define separate rules for each type of action that requires different groups.
- Interfaces—The interface or interface role for which you are configuring the rule.

Click **OK** when you are finished defining your rule.

Step 4 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.

Step 5 Select **Firewall > Settings > AAA Firewall** (in Device or Policy view) to open the AAA Firewall Settings Page, Advanced Setting Tab, page 15-20. Configure the AAA firewall settings:

- If you configured rules for HTTP authentication, you should select **Use Secure HTTP Authentication**. This ensures that the username and password entered for HTTP authentication are encrypted. If you do not select this option, the credentials are sent in clear text, which is insecure.

Tip If you select this option, ensure that you do not configure 0 for the user authentication timeout (**timeout uauth 0**, configured in the **Platform > Security > Timeouts** policy), or users might be repeatedly prompted for authentication, making the feature disruptive to your network.

- If you configured authentication for HTTP or HTTPS traffic on an interface, you should consider adding the interface to the Interactive Authentication table. When you enable an interface for interactive authentication, users get an improved authorization web page, one that is the same for both HTTP and HTTPS.

  Click **Add Row** to add the interface to the table. Select whether the interface should listen for HTTP or HTTPS traffic (add the interface twice to listen for both protocols), and the port to listen on if not the default port for the protocol (80 and 443, respectively). Select **Redirect network users for authentication request** so that network access traffic gets the improved authentication prompt; if you do not select the option, only users trying to log into the device get the prompt.

Note You might want to continue to use basic HTTP authentication if: you do not want the security appliance to open listening ports; if you use NAT on a router and you do not want to create a translation rule for the web page served by the security appliance; basic HTTP authentication might work better with your network. For example non-browser applications, like when a URL is embedded in email, might be more compatible with basic authentication.

- For FWSM devices, you can also disable the authentication challenge for protocols you have otherwise configured to require authentication. You can also add interfaces to the Clear Connections table to ensure that active connections for users whose authentication has timed out are cleared and do not hang.
- If you want to exempt some devices from your AAA rules based on their media access control (MAC) address, click the **MAC Exempt List** tab to open the AAA Firewall Page, MAC-Exempt List Tab, page 15-26. Enter a name for the exemption list, and then click the **Add Row** button and fill in the Firewall AAA MAC Exempt Setting Dialog Box, page 15-27 to add the MAC address to the table with a permit rule. You might want to do this for secure, trusted devices.

  The order of entries matters, so ensure that any specific entries that are covered by broader entries come before the broad entries in the table. The device processes the list in order and the first match is applied to the host. For more detailed information about how the entries on the MAC exempt list are processed, see AAA Firewall Page, MAC-Exempt List Tab, page 15-26.

Step 6 If you are configuring authentication rules using a RADIUS server, and you include per-user ACL configurations in the user profiles, enable per-user downloadable ACLs on the interface. (RADIUS authentication automatically includes authorization checking.) For information on configuring per-user ACLs, see the information on configuring RADIUS authorization in the *Cisco ASA 5500 Series Configuration Guide Using the CLI* at http://www.cisco.com/en/US/docs/security/asa/asa82/configuration/guide/access_fwaaa.html.

a. Select **Firewall > Settings > Access Control** (in Device or Policy view) to open the Access Control Settings Page, page 16-24.

b. Click the Add Row button beneath the interface table and fill in the Firewall ACL Setting Dialog Box, page 16-26 with at least these options:

   - Enter the interface or interface role on which you are performing authorization.
   - Select **Per User Downloadable ACLs**.

c. Click **OK** to save your changes.

Configuring AAA Rules for IOS Devices

When you configure AAA rules for an IOS device, you are configuring authentication proxy (AuthProxy) admission control policies. These policies define who is allowed to make HTTP, HTTPS, FTP, and Telnet connections through (not to) the device. To fully configure authentication proxy, you must configure several policies, not just the AAA rules policy.

The following procedure covers all policies you would need to configure to supply full authentication, authorization, and accounting support for authorization proxy. You do not need to configure options for features you do not need.

Related Topics

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Creating a New Shared Policy, page 5-54
- Modifying Policy Assignments in Policy View, page 5-54
- Understanding Networks/Hosts Objects, page 6-80

- Understanding Interface Role Objects, page 6-73
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Understanding AAA Server and Server Group Objects, page 6-27
- Understanding Interface Role Objects, page 6-73

Step 1 Do one of the following to open the AAA Rules Page, page 15-10:

- (Device view) Select **Firewall > AAA Rules** from the Policy selector.
- (Policy view) Select **Firewall > AAA Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add and Edit AAA Rule Dialog Boxes, page 15-13.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10.

Step 3 Configure the rule. Following are the highlights of what you typically need to decide. For specific information on configuring the fields, see Add and Edit AAA Rule Dialog Boxes, page 15-13.

- **Authentication Action**—Select this option. Authentication rules are the only type of rule you can configure in the AAA rules policy for IOS devices.
- Permit or Deny—Whether you are subjecting the identified traffic to AAA control (permit) or you are exempting it from AAA control (deny). Any denied traffic is not prompted for authentication and is allowed to pass unauthenticated, although your access rules might drop the traffic.
- Source and Destination addresses—If the rule should apply no matter which addresses generated the traffic or their destinations, use “All-Addresses” as the source or destination. If the rule is specific to a host or network, enter the addresses or network/host objects. For information on the accepted address formats, see Specifying IP Addresses During Policy Definition, page 6-87.
- Source and Destination Security Groups (ASA 9.0+ only)—You can specify TrustSec security groups used to filter traffic in addition to the source and destination addresses. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups.
- Source Users (ASA 8.4.2+ only)—You can further define the traffic source by specifying Active Directory (AD) user names (in the format NetBIOS_DOMAIN\username), user groups (NetBIOS_DOMAIN\\user_group), or identity user group objects that define the names and groups. The user specification is conjoined to the source address to limit the match to user addresses within the source address range. For more information, see Configuring Identity-Based Firewall Rules, page 13-21 and Creating Identity User Group Objects, page 13-19.
- Services—You can specify any type of service for authentication and authorization rules; however, the user is prompted to authenticate only for HTTP, HTTPS, FTP, and Telnet connections. Thus, if you specify something other than these services, the user must first attempt one of these connections and successfully authenticate (and be authorized, if you include that action) before any other types of connections are allowed. For accounting rules, you can specify any TCP or UDP service (or simply TCP and UDP themselves), if you want to account for all types of traffic.
- Interfaces—The interface or interface role for which you are configuring the rule.

- Service triggering the authentication proxy—Select the checkboxes for the type of traffic you want to trigger user authentication: HTTP, FTP, or Telnet. You can select any combination. If you want to trigger the proxy for HTTPS support, select HTTP and perform the HTTPS configuration that is explained in a subsequent step in this procedure.

Click **OK** when you are finished defining your rule.

Step 4 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.

Step 5 Select **Firewall > Settings > AuthProxy** (in Device or Policy view) to open the AAA Page, page 15-28. Configure the authentication proxy settings:

- **Authorization server groups**—If you want all of your authentication rules to also perform user authorization, specify the list of AAA server group policy objects that identify the TACACS+ or RADIUS servers that control authorization. You can also specify LOCAL to use the user database defined on the device. If you do not specify a server group, authorization is not performed.

Tip You must configure per-user ACLs in your AAA server to define the privileges you want to apply to each user. When configuring authorization, specify **auth-proxy** as the service (e.g. service = auth-proxy), with a privilege level of 15. For more information on configuring the AAA server, including information on configuring authentication proxy in general, see the "Configuring the Authentication Proxy" section in the *Cisco IOS Security Configuration Guide: Securing User Services, Release 12.4T* at http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_cfg_authen_prxy_ps6441_TSD_Products_Configuration_Guide_Chapter.html.

- **Accounting server groups**—If you want to perform accounting for all of your authentication rules, specify the list of AAA server group policy objects that identify the TACACS+ or RADIUS servers that perform accounting. If you do not specify a server group, no accounting is performed. When performing accounting, also configure the following options as appropriate:
  - If you specify more than one server group, consider selecting **Use Broadcast for Accounting**. This option sends accounting records to the primary server in each server group.
  - The **Accounting Notice** option defines when the server is notified. The default is to notify the server at the start and stop of a connection, but you can select to only send stop notices (or none at all).
- You can also customize authentication banners for each service, and on the Timeout tab, you can change the default idle and absolute session timeouts globally or for each interface.

Step 6 Select **Platform > Device Admin > AAA** (in policy view, this is in the Router Platform folder) to open the AAA Policy Page, page 63-6. Configure these options on the Authentication tab:

- Select **Enable Device Login Authentication**.
- Enter the list of server groups that will control authentication in priority order. Typically, you will use at least some of the same LDAP, RADIUS, or TACACS+ server groups used in the AuthProxy policy. However, this policy also defines device login control, so you might want to include some other server groups. For more information, see AAA Page—Authentication Tab, page 63-6.

Step 7 If you are using the authentication proxy with HTTP connections, and you also want to use the proxy with HTTPS connections, select **Platform > Device Admin > Device Access > HTTP** (in policy view, this is in the Router Platform folder) to open the HTTP Policy Page, page 63-31. Configure these options:

- Select **Enable HTTP** and **Enable SSL** if they are not already selected.

- On the AAA tab, ensure that the configuration for login access to the device is appropriate. If you are using AAA to control access through the device, you might want to use it for access to the device.

AAA Rules Page

Use the AAA Rules page to configure AAA rules for device interfaces. AAA rules configure network access control (called authentication proxy on IOS devices), which forces the user to authenticate when attempting network connections that traverse the device. Authenticated traffic can also be required to undergo authorization (where after the user enters a valid user name and password, the AAA server is checked to verify that the user is authorized for network access). You can also configure accounting rules, even for unauthenticated traffic, to provide information you can use for billing, security, and resource allocation purposes.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate policies and objects for configuring IPv4 and IPv6 AAA rules were "unified," meaning one set of AAA rules in which you can use either IPv4 or IPv6 addresses, or a mixture of both. (See Policy Object Changes in Security Manager 4.4, page 1-10 for additional information.) In Policy view, IPv4 and unified versions of the AAA policy type are provided. In addition, a utility that you can use to convert existing IPv4 policies is provided (see Converting IPv4 Rules to Unified Rules, page 12-28). The following descriptions apply to all versions of the AAA rule table, except where noted.

If you assign an IPv4 AAA-rule shared policy to a 9.0+ device, you will no longer be able to assign unified versions of those policies to that device. Likewise, if you assign a unified AAA-rule shared policy to a 9.0+ device, you will no longer be able to assign IPv4 versions of those shared policies to that device--the device will not be included in the list of available devices on the Assignments tab for the shared policy.

AAA rule configuration is complex and differs significantly based on the operating system. Carefully read the following topics before configuring AAA rules:

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4
- Configuring AAA Rules for IOS Devices, page 15-7

Tip Disabled rules are shown with hash marks covering the table row. When you deploy the configuration, disabled rules are removed from the device. For more information, see Enabling and Disabling Rules, page 12-20.

Navigation Path

To access the AAA Rules page, do one of the following:

- (Device view) Select a device, then select **Firewall > AAA Rules** from the Policy selector.
- (Policy view) Select **Firewall > AAA Rules** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Policies > AAA Rules**.

Related Topics.

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Moving Rules and the Importance of Rule Order, page 12-19
- Using Sections to Organize Rules Tables, page 12-20
- Using Rules Tables, page 12-8
- Filtering Tables, page 1-48

Field Reference

Table 15-1 AAA Rules Page

Element	Description
Expand all rows/Collapse all rows	Use these buttons to expand or collapse all sections in the rules table. **Note** The buttons are located in the upper-right corner of the Filter area above the access rules table.
Conflict Indicator icons	Identifies conflicts and provides a quick visual representation of the type of conflict. For more details, including types of conflicts and the actions you can take from this column, see Understanding the Automatic Conflict Detection User Interface, page 16-30.
No.	The ordered rule number.
Permit	Whether the defined traffic will be subject to the rule (Permit) or exempted from the rule (Deny): • Permit—Shown as a green check mark. • Deny—Shown as a red circle with slash.
Sources	The sources of traffic for this rule; can be networks, security groups (ASA 9.0+ only), and users. Multiple entries are displayed on separate lines within the table cell.
Destinations	The destinations for this rule; can be networks and security groups (ASA 9.0+ only). Multiple entries are displayed on separate lines within the table cell.
Service	The services or service objects that specify the protocol and port of the traffic to which the rule applies. Multiple entries are displayed on separate lines within the table cell. See Understanding and Specifying Services and Service and Port List Objects, page 6-100.
Interface	The interfaces or interface roles to which the rule is assigned. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. Multiple entries are displayed as separate subfields within the table cell. See Understanding Interface Role Objects, page 6-73.

Table 15-1 AAA Rules Page (continued)

Element	Description
Action	The type of AAA control defined by this rule: • Authenticate—Users making connections through the device must authenticate with their username and password. Protocols requiring authentication are defined by the Service field (for ASA/PIX/FWSM devices) or the AuthProxy methods (for IOS devices). • Authorize—Authenticated users are also checked with the AAA server to ensure that they are authorized to make the connection (ASA/PIX/FWSM only). • Account—Accounting records for the identified traffic are sent to the AAA server (ASA/PIX/FWSM only). You can right-click the Action cell in an existing AAA rule and choose Edit Action to change your selections. See Edit AAA Option Dialog Box, page 15-19 for more information.
AAA Method (IOS) (not presented for ASA 9.0+ devices)	The authentication method for this rule: Web Authorization Proxy (Auth-Proxy), HTTP Basic, or Windows NT LAN Manager (NTLM)
AuthProxy	The protocols that require authentication using the authentication proxy method. This applies only to IOS devices. You can right-click the AuthProxy cell in an existing AAA rule and choose Edit AuthProxy to change your selections. See AuthProxy Dialog Box, page 15-19 for more information.
Server Group	The AAA server group that provides the authentication, authorization, or accounting support defined in the rule. This group is used for ASA/PIX/FWSM devices only. For information on configuring AAA servers for IOS devices for use with these rules, see Configuring AAA Rules for IOS Devices, page 15-7. You can right-click the Server Group cell in an existing AAA rule and choose Edit Server Group to change your selections. See Edit Server Group Dialog Box, page 15-19 for more information.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of the rule, if any.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Page elements below the rules table	
Query	Click this button to run a policy query, which can help you evaluate your rules and identify ineffective rules. See Generating Policy Query Reports, page 12-28

Table 15-1 AAA Rules Page (continued)

Element	Description
Find and Replace button (binoculars icon)	Click this button to search for various types of items within the table and to optionally replace them. See Finding and Replacing Items in Rules Tables, page 12-16.
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope or section. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Add Row button	Click this button to add a rule to the table after the selected row using the Add and Edit AAA Rule Dialog Boxes, page 15-13. If you do not select a row, the rule is added at the end of the local scope. For more information about adding rules, see Adding and Removing Rules, page 12-9.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells. For more information, see Editing Rules, page 12-10.
Delete Row button	Click this button to delete the selected rule.

Right-click Menu

A right-click menu is also available. This menu provides access to many of the functions listed above; the options presented depend on the location right-clicked:

- If you right-click a rule in the table, the options may include editing functions relative to the specific table cell right-clicked. For example, the command “Edit Server Group” is included when you right-click a Server Group cell. See Editing Rules, page 12-10 for more information.
- The Combine Rules option is also included in the right-click menu. See Combining Rules, page 12-22 for more information.

Add and Edit AAA Rule Dialog Boxes

Use the Add and Edit AAA Rules dialog boxes to add and edit AAA rules. AAA rule configuration is more complex than just filling in this dialog box, and differs significantly based on the operating system. Carefully read the following topics before configuring AAA rules:

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4
- Configuring AAA Rules for IOS Devices, page 15-7

Navigation Path

From the AAA Rules Page, page 15-10, click the **Add Row** button or select a row and click the **Edit Row** button.

Related Topics

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10

Field Reference

Table 15-2 Add and Edit AAA Rules Dialog Boxes

Element	Description
Enable Rule	Whether to enable the rule, which means the rule becomes active when you deploy the configuration to the device. Disabled rules are shown overlain with hash marks in the rule table. For more information, see Enabling and Disabling Rules, page 12-20.
Action (Permit/Deny)	Whether the defined traffic will be subject to the rule (Permit) or exempted from the rule (Deny). For example, if you create an authentication deny rule for the 10.100.10.0/24 network to any destination using the HTTP service, users on this network are not prompted to authenticate with the device when making HTTP requests.

Table 15-2 Add and Edit AAA Rules Dialog Boxes (continued)

Element	Description
Sources	Provide traffic sources for this rule; can be networks, security groups, and users. You can enter values or object names, or Select objects, for one or more of the following types of sources: • Network – You can specify a various network, host and interface definitions, either individually or as objects. If you Select an interface object as a source, the dialog box displays tabs to differentiate between hosts/networks and interfaces. The "All-Address" objects do not restrict the rule to specific hosts, networks, or interfaces. These addresses are IPv4 or IPv6 addresses for hosts or networks, network/host objects, interfaces, or interface roles. **Note** You can only specify a fully qualified domain name (FQDN) by providing an FQDN network/host object, or a group object that includes an FQDN object. You cannot directly type in an FQDN. See Understanding Networks/Hosts Objects, page 6-80, Specifying IP Addresses During Policy Definition, page 6-87 and Understanding Interface Role Objects, page 6-73 for additional information about these definitions. • Security Groups (ASA 9.0+) – Enter or Select the name or tag number for one or more source security groups for the rule, if any. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. • Users – Enter or Select the Active Directory (AD) user names, user groups, or identity user group objects for the rule, if any. You can enter any combination of the following: – Individual user names: NetBIOS_DOMAIN\username – User groups (note the double \): NetBIOS_DOMAIN\\user_group – Identity user group object names. For more information, see: – Selecting Identity Users in Policies, page 13-21 – Configuring Identity-Based Firewall Rules, page 13-21 – Creating Identity User Group Objects, page 13-19 **Note** Enter more than one value in any of these fields by separating the items with commas. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.

Table 15-2 Add and Edit AAA Rules Dialog Boxes (continued)

Element	Description
Destinations	Provide traffic destinations for this rule; can be networks or security groups. As with Sources, you can enter values or object names, or Select objects, for one or more destinations of Network and Security Group (ASA 9.0+) type.
Services	The services that define the type of traffic upon which to act. You can enter or Select any combination of service objects and service types (which are typically a protocol and port combination). Enter more than one value by separating the items with commas. It is important that you select the service type carefully based on the device type: • For IOS devices, only the protocols you select with the authorization proxy check boxes at the bottom of the dialog box are used for AAA control, so you can use IP as the protocol. • For ASA, PIX, and FWSM devices, although you can force authentication for any type of traffic, the security appliance prompts only for HTTP/HTTPS, FTP, and Telnet traffic. If you specify a service other than one of these, users are prevented from making any connection through the appliance until they try one of these services and successfully authenticate. If the rule is only for accounting, you can specify any TCP or UDP protocols for which you want to create records. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100. **Note** Due to an issue in PIX 6.3 and FWSM devices, if you specify a service with a source port, no traffic is authenticated. Therefore, source ports are ignored when the CLI is generated from your rule for these device types.
Interface	The interface or interface role object that identifies the interface from which to authenticate, authorize, or account users. Enter the name of the interface or interface role, or click **Select** to select it from a list or to create a new interface role object. For authentication rules on ASA and PIX devices, you can modify how this interface authenticates HTTP/HTTPS traffic by using the **Firewall > Settings > AAA Firewall** policy. Configuring the interface as an HTTP/HTTPS listening port can improve the authentication experience for users. For more information, see Understanding How Users Authenticate, page 15-2 and AAA Firewall Settings Page, Advanced Setting Tab, page 15-20.
Description	An optional description of the rule (up to 1024 characters).

Table 15-2 Add and Edit AAA Rules Dialog Boxes (continued)

Element	Description
The Authentication Action, Authorization Action, and Accounting Action check boxes define the types of rules that will be generated on the device. Each type generates a separate set of commands, but if you select more than one option, your other selections in this dialog box are limited to those supported by all selected actions. You can right-click the Action cell in an existing AAA rule and choose Edit Action to change your selections. See Edit AAA Option Dialog Box, page 15-19 for more information.	
Authentication Action User-Identity	• Authentication—Users must supply a user name and password to make a connection through the device. For ASA, PIX, and FWSM devices, what you enter in the Services field determines which protocols require authentication, although the device will prompt only for HTTP, HTTPS, FTP, and Telnet connections. For IOS devices, the protocols that require authentication are based on the authorization proxy check boxes you select at the bottom of the dialog box. – User-Identity (ASA 8.4(2+) only.)—For ASA devices, when you select Authentication Action, you also have the option to select User-Identity. This option indicates that the device should use the identity-firewall domain mappings defined in the Identity Options policy to authenticate users instead of the AAA Server Group setting in the AAA rule. If the user enters a domain name, the AD server associated with the domain is queried. Otherwise, the AD server associated with the default domain is queried. See Identifying Active Directory Servers and Agents, page 13-8.
Authorization Action (PIX/ASA/FWSM)	Authorization—After successful authentication, the AAA server is also checked to determine if the user is authorized to make the requested connection. If you specify a RADIUS server for authentication rules, authorization happens without you having to configure authorization rules. If you are using a TACACS+ server, you must create separate authorization rules.
Accounting Action (PIX/ASA/FWSM)	Accounting—Accounting records will be sent to the TACACS+ or RADIUS server for the TCP and UDP protocols specified in the Services field. If you also configure authentication, these records are per-user; otherwise, they are based on IP address. For IOS devices, accounting is configured in the **Firewall > Settings >ScanSafe Web Security** policy, not in AAA rules, and applies only to the protocols you select for authentication proxy.

Table 15-2 Add and Edit AAA Rules Dialog Boxes (continued)

Element	Description
AAA Server Group (PIX, ASA, FWSM)	The AAA server group policy object that defines the AAA server that should provide authentication, authorization, or accounting for the traffic defined in the rule. Enter the name of the policy object or click **Select** to select it from a list or to create a new object. You must select a type of server that can perform all actions defined in the rule. For example, the local database (defined on the device) cannot provide authorization services. If you use a RADIUS server for authentication, it automatically provides authorization services, but you cannot define an authorization rule that uses a RADIUS server. You can use a mix of server groups for different actions for the same source/destination pair by creating separate rules for the desired combination of authentication, authorization, and accounting actions. For more information on AAA server group objects, see Understanding AAA Server and Server Group Objects, page 6-27. **Tips** • If you select Authenticate Action and User-Identity, but not the Authorization or Accounting actions, any server you specify here is ignored. Do not select a server to avoid validation warnings. • AAA server groups for IOS devices are defined in other policies. For a complete explanation of the configuration, see Configuring AAA Rules for IOS Devices, page 15-7. • You can right-click the Server Group cell in an existing AAA rule and choose Edit Server Group to change your selections. See Edit Server Group Dialog Box, page 15-19 for more information.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Method (IOS) (not presented for ASA 9.0+ devices)	Choose Auth-Proxy, HTTP-basic, or NTLM. If you choose Auth-Proxy, the following options are available: • HTTP • FTP • Telnet Specify the protocols for which you want to enforce authentication using the authentication proxy. If you select HTTP, you can also configure HTTPS authentication proxy by enabling SSL on the device. For specific information, see Configuring AAA Rules for IOS Devices, page 15-7. You can right-click the AuthProxy cell in an existing AAA rule and choose Edit AuthProxy to change your selections. See AuthProxy Dialog Box, page 15-19 for more information.

Edit AAA Option Dialog Box

Use the Edit AAA Option dialog box to select whether the rule performs authentication (with or without user identity), authorization, or accounting. Authorization and accounting rules work only on ASA, PIX, and FWSM devices. For a complete explanation of these options, see the related explanations in the following topics:

- Add and Edit AAA Rule Dialog Boxes, page 15-13
- Understanding AAA Rules, page 15-1

Navigation Path

Right-click the Action cell in a AAA rule (on the AAA Rules Page, page 15-10) and select **Edit AAA**.

AuthProxy Dialog Box

Use the AuthProxy dialog box to edit the authorization proxy settings in a AAA rule. For IOS devices, select the protocols (HTTP, FTP, or Telnet) for which you want to enforce authentication using the authentication proxy. If you select HTTP, you can also configure HTTPS authentication proxy by enabling SSL on the device. For specific information, see Configuring AAA Rules for IOS Devices, page 15-7.

Navigation Path

Right-click the AuthProxy cell in a AAA rule (on the AAA Rules Page, page 15-10) and select **Edit AuthProxy**.

Edit Server Group Dialog Box

Use the Edit Server Group dialog box to edit the AAA server group used in a AAA rule, which defines the AAA server that should provide authentication, authorization, or accounting for the traffic defined in the rule. Enter the name of the policy object or click **Select** to select it from a list or to create a new object. For more information on AAA server group objects, see Understanding AAA Server and Server Group Objects, page 6-27.

You must select a type of server that can perform all actions defined in the rule. For example, the local database (defined on the device) cannot provide authorization services. If you use a RADIUS server for authentication, it automatically provides authorization services, but you cannot define an authorization rule that uses a RADIUS server. Unlike the Add and Edit AAA Rule Dialog Boxes, page 15-13, this dialog box does not validate your selection.

Note This setting applies only to ASA, PIX, and FWSM devices. AAA server groups for IOS devices are defined in other policies. For a complete explanation of the configuration, see Configuring AAA Rules for IOS Devices, page 15-7.

Navigation Path

Right-click the Server Group cell in a AAA rule (on the AAA Rules Page, page 15-10) and select **Edit Server Group**.

AAA Firewall Settings Policies

The AAA firewall settings policy configurations influence the behavior of your AAA rules.

This section contains the following topics:

- AAA Firewall Settings Page, Advanced Setting Tab, page 15-20
- AAA Firewall Page, MAC-Exempt List Tab, page 15-26
- AAA Page, page 15-28

AAA Firewall Settings Page, Advanced Setting Tab

Use the AAA Firewall settings policy to configure optional settings to refine how your AAA rules policy behaves. This topic describes the settings available on the Advanced Setting tab; for information on the MAC Exempt List tab, see AAA Firewall Page, MAC-Exempt List Tab, page 15-26.

Navigation Path

To access the AAA Firewall settings page, do one of the following:

- (Device view) Select an ASA, PIX, or FWSM device, select **Firewall > Settings > AAA Firewall**; select the **Advanced Setting** tab if necessary.
- (Policy view) Select **Firewall > Settings > AAA Firewall** from the Policy Type selector. Create a new policy or select and existing one, then select the **Advanced Setting** tab if necessary.
- (Map view) Right-click an ASA, PIX, or FWSM device and select **Edit Firewall Settings > AAA Firewall**, then select the **Advanced Setting** tab if necessary.

Related Topics

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4

Field Reference

Table 15-3 Advanced Setting Tab, AAA Firewall Settings Page

Element	Description
Use Secure HTTP Authentication	Whether to require users making HTTP requests that traverse the security appliance to first authenticate with the security appliance using SSL (HTTPS). The user is prompted for username and password. Secure HTTP authentication offers a secure method for user authentication to the security appliance prior to allowing HTTP-based web requests to traverse the security appliance. This is also called HTTP cut-through proxy authentication. If you select this option, ensure that your access rules do not block HTTPS traffic (port 443), and that any PAT configuration also includes port 443. Also, be aware that a maximum of 16 concurrent authentications are allowed, and that if you configure 0 for the user authentication timeout (**timeout uauth 0**, configured in the **Platform > Security > Timeouts** policy) users might be repeatedly prompted for authentication, making the feature disruptive to your network. **Tip** If you do not select this option, HTTP authentication sends the username and password in clear text.
Enable Proxy Limit Maximum Concurrent Proxy Limit per User	Whether to allow proxy connections. If you enable proxies, you must set a limit on the number of proxy connections allowed for each user, from 1 to 128. The device default is 16, but you must specify a number.

Table 15-3 Advanced Setting Tab, AAA Firewall Settings Page (continued)

Element	Description
Enable Virtual HTTP	Whether to configure a virtual HTTP server. This feature redirects all HTTP connections that require AAA authentication to the virtual HTTP server on the ASA. The ASA prompts for the AAA server username and password. After the AAA server authenticates the user, the ASA redirects the HTTP connection back to the original server, but it does not include the AAA server username and password. Because the username and password are not included in the HTTP packet, the HTTP server prompts the user separately for the HTTP server username and password. For more information, see Understanding How Users Authenticate, page 15-2. For inbound users (from lower security to higher security), you must also include the virtual HTTP address as a destination interface in the access rule applied to the source interface. Moreover, you must add a static NAT rule for the virtual HTTP IP address, even if NAT is not required. An identity NAT rule is typically used (where you translate the address to itself). For outbound users, there is an explicit permit for traffic, but if you apply an access rule to an inside interface, be sure to allow access to the virtual HTTP address. A static NAT rule is not required. To configure a virtual HTTP server: 1. Select the Enable Virtual HTTP check box. 2. Enter the IP address or select a Networks/Hosts object representing the virtual HTTP server. Make sure this address is an unused address that is routed to the ASA. For example, if you perform NAT for inside addresses accessing an outside server, and you want to provide outside access to the virtual HTTP server, you can use one of the global NAT addresses for the virtual HTTP server address. 3. (Optional) If you are using text-based browsers, where redirection does not happen automatically, select the Warning check box. This enables an alert to notify users when the HTTP connection is being redirected.

Table 15-3 Advanced Setting Tab, AAA Firewall Settings Page (continued)

Element	Description
Enable Virtual Telnet	Whether to configure a virtual Telnet server. When an unauthenticated user connects to the virtual Telnet IP address, the user is challenged for a username and password, and then authenticated by the AAA server. After the user is authenticated, the message "Authentication Successful" appears. Then the user can successfully access other services that require authentication. For inbound users (from lower security to higher security), you must also include the virtual Telnet address as a destination interface in the access rule applied to the source interface. In addition, you must add a static NAT rule for the virtual Telnet IP address, even if NAT is not required. An identity NAT rule is typically used (where you translate the address to itself). For outbound users, there is an explicit permit for traffic, but if you apply an access rule to an inside interface, be sure to allow access to the virtual Telnet address. A static NAT rule is not required. To configure a virtual Telnet server: 1. Select the Enable Virtual Telnet check box. 2. Enter the IP address or select a Networks/Hosts object representing the virtual Telnet server. Make sure this address is an unused address that is routed to the ASA. For example, if you perform NAT for inside addresses accessing an outside server, and you want to provide outside access to the virtual Telnet server, you can use one of the global NAT addresses for the virtual Telnet server address.
Interactive Authentication table (ASA/PIX 7.2.2+)	Use this table to identify the interfaces that should listen for HTTP or HTTPS traffic for authentication. If your AAA rules require authentication for these protocols on interfaces designated in this table, the user is presented with an improved authentication web page as opposed to the default authentication pages used by the appliance. These pages are also used for authenticating direct connections to the device. • To add an interface to the table, click the **Add Row** button and fill in the Interactive Authentication Configuration Dialog Box, page 15-24. • To edit a setting, select it and click the **Edit Row** button. • To delete a setting, select it and click the **Delete Row** button.

Table 15-3 Advanced Setting Tab, AAA Firewall Settings Page (continued)

Element	Description
Disable FTP Authentication Challenge Disable HTTP Authentication Challenge Disable HTTPS Authentication Challenge Disable Telnet Authentication Challenge (All options FWSM 3.x+ only.)	Whether to disable authentication challenges for the indicated protocols. By default, the FWSM prompts the user for a username and password when a AAA rule enforces authentication for traffic in a new session and the protocol of the traffic is FTP, Telnet, HTTP, or HTTPS. In some cases, you might want to disable the authentication challenge for one or more of these protocols. If you disable challenge authentication for a particular protocol, traffic using that protocol is allowed only if the traffic belongs to a session previously authenticated. This authentication can be accomplished by traffic using a protocol whose authentication challenge remains enabled. For example, if you disable challenge authentication for FTP, the FWSM denies a new session using FTP if the traffic is included in an authentication AAA rule. If the user establishes the session with a protocol whose authentication challenge is enabled (such as HTTP), FTP traffic is allowed.
Clear Connections When Uauth Timer Expires table (FWSM 3.2+ only.)	Use this table to identify the interfaces and source addresses where you want to force any active connections to close immediately after the user authentication times out or when you clear the authentication session with the **clear uauth** command. (User authentication timeouts are defined in the **Platform > Security > Timeouts** policy.) For any interface/source address pairs not listed in this table, active connections are not terminated even though the user authentication session expired. • To add any interface and source address pair, click the **Add Row** button and fill in the Clear Connection Configuration Dialog Box, page 15-25. • To edit a setting, select it and click the **Edit Row** button. • To delete a setting, select it and click the **Delete Row** button.

Interactive Authentication Configuration Dialog Box

Use the Interactive Authentication Configuration dialog box to configure an interface to listen for HTTP or HTTPS traffic to authenticate network users. The authentication web page used by a listening port provides an improved user experience compared to the default authentication pages used for these protocols. The authentication pages are used for connections directly to the device and if you select the redirection option, also for through traffic if your AAA rules policy requires HTTP/HTTPS network access authentication. For more information, see Understanding How Users Authenticate, page 15-2.

Navigation Path

Go to the AAA Firewall Settings Page, Advanced Setting Tab, page 15-20 and click the **Add Row** button beneath the Interactive Authentication table, or select an item in the table and click the **Edit Row** button.

Related Topics

- Understanding AAA Rules, page 15-1
- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4

Field Reference

Table 15-4 Interactive Authentication Configuration Dialog Box

Element	Description
Protocol	The protocol that you want to listen for, either HTTP or HTTPS. If you want to listen for both protocols on an interface, add the interface to the table twice.
Interface	The interface or interface role on which to enable listeners. Enter the name of the interface or interface role, or click **Select** to select it from a list or to create a new interface role.
Port	The port number that the security appliance listens on for this protocol if other than the default, which is 80 (HTTP) and 443 (HTTPS).
Redirect network users for authentication request	Whether to redirect users who are making requests through the device to the authentication web page served by the security appliance. If you do not select this option, only traffic directed to the interface is prompted with the improved authentication web page.

Clear Connection Configuration Dialog Box

Use the Clear Connection Configuration dialog box to identify the source addresses whose active connections to close immediately after the user authentication times out or when you clear the authentication session with the **clear uauth** command. You must specify the interfaces on which those sessions should be cleared. These settings are used only for FWSM 3.2+ devices.

User authentication timeouts are defined in the **Platform > Security > Timeouts** policy.

Navigation Path

Go to the AAA Firewall Settings Page, Advanced Setting Tab, page 15-20 and click the **Add Row** button beneath the Clear Connections When Uauth Timer Expires table, or select an item in the table and click the **Edit Row** button.

Field Reference

Table 15-5 Clear Connection Configuration Dialog Box

Element	Description
Interface	The interfaces or interface roles for which you are configuring settings. Enter the name or click **Select** to select the interface or interface role or to create a new role. Separate multiple entries with commas.
Source IP Address/Netmask	The host or network addresses for which you want to clear connections immediately when the user authentication timer expires. The list can include host IP addresses, network addresses, address ranges, or network/host objects Separate multiple addresses with commas. For more information on entering addresses, see Specifying IP Addresses During Policy Definition, page 6-87.

AAA Firewall Page, MAC-Exempt List Tab

Use the MAC Exempt List tab of the AAA Firewall settings policy to identify hosts that should be exempt from authentication and authorization for ASA, PIX, and FWSM 3.x+ devices. For example, if the security appliance authenticates TCP traffic originating on a particular network but you want to allow unauthenticated TCP connections from a specific server, create a rule permitting traffic from the MAC address of the server.

You can use masks to create rules for groups of MAC addresses. For example, if you want to exempt all Cisco IP phones whose MAC addresses start with 0003.e3, create a permit rule for 0003.e300.0000 with the mask ffff.ff00.0000. (An f in a mask exactly matches the corresponding number in the address, whereas a 0 matches anything.)

Deny rules are necessary only if you are permitting a group of MAC addresses but there are some addresses within the permitted group that you want to require to use authentication and authorization. Deny rules do not prohibit traffic; they simply require the host to go through normal authentication and authorization. For example, if you want to allow all hosts with MAC addresses that start with 00a0.c95d, but you want to force 00a0.c95d.0282 to use authentication and authorization, enter these rules in order:

1. Deny 00a0.c95d.0282 ffff.ffff.ffff
2. Permit 00a0.c95d.0000 ffff.ffff.0000

When you deploy the policy to the device, these entries are configured using the **mac-list** and **aaa mac-exempt** commands.

Tip The MAC exempt list is processed on a first match basis. Thus, the order of entries matters. If you want to permit a group of MAC addresses, but deny a subset of them, the deny rule must come before the permit rule. However, Security Manager does not allow you to order MAC exempt rules: they are implemented in the order shown. If you sort the table, your policy changes. If your entries do not depend on each other, this does not matter. Otherwise, ensure that you enter rows in the proper order.

Navigation Path

To access the MAC Exempt List tab, do one of the following:

- (Device view) Select an ASA, PIX, or FWSM device, then select **Firewall > Settings > AAA Firewall**. Select the **MAC-Exempt List** tab.
- (Policy view) Select **Firewall > Settings > AAA Firewall** from the Policy Type selector. Create a new policy or select an existing one, then select the **MAC-Exempt List** tab.
- (Map view) Right-click an ASA, PIX, or FWSM device and select **Edit Firewall Settings > AAA Firewall**, then select the **MAC-Exempt List** tab.

Related Topics

- Configuring AAA Rules for ASA, PIX, and FWSM Devices, page 15-4
- Filtering Tables, page 1-48

Field Reference

Table 15-6 MAC-Exempt List Tab, AAA Firewall Settings Page

Element	Description
MAC-Exempt List Name	The name of the MAC exempt list.

Table 15-6 MAC-Exempt List Tab, AAA Firewall Settings Page (continued)

Element	Description
MAC Exempt List table	The MAC exempt rules that you want to implement. The table shows the MAC addresses and masks (in hexadecimal) and whether you are permitting them (exempting them from authentication and authorization) or denying them (making them go through standard authentication and authorization). The device processes the entries in order and uses the first match (not the best match). • To add an exemption rule, click the **Add Row** button and fill in the Firewall AAA MAC Exempt Setting Dialog Box, page 15-27. • To edit an exemption rule, select it and click the **Edit Row** button. • To delete an exemption rule, select it and click the **Delete Row** button.

Firewall AAA MAC Exempt Setting Dialog Box

Use the Firewall AAA MAC Exempt Setting dialog box to add and edit exemption entries in the MAC Exempt List table. The security appliance skips authentication and authorization for hosts associated with permitted MAC addresses.

Navigation Path

Go to the AAA Firewall Page, MAC-Exempt List Tab, page 15-26 and click the **Add Row** button beneath the MAC Exempt List table, or select an item in the table and click the **Edit Row** button.

Field Reference

Table 15-7 Firewall AAA MAC Exempt Setting Dialog Box

Element	Description
Action	The action you want to take for the hosts that use the specified MAC addresses: • Permit—Exempts the host from authentication and authorization. • Deny—Forces the host to go through authentication and authorization.
MAC Address	The MAC address of the hosts in standard 12-digit hexadecimal format, such as 00a0.cp5d.0282. You can enter complete MAC addresses or partial addresses. For partial addresses, you can enter 0 for digits you are not matching.
MAC Mask	The mask to apply to the MAC address. Use f to match a digit exactly, 0 to match any digit at that place: • To specify an exact match of the address, enter ffff.ffff.ffff. • To match an address pattern, enter 0 for any digit for which you want to match any character. For example, ffff.ffff.0000 matches all addresses that have the same first 8 digits.

AAA Page

Use the AAA firewall settings policy to identify the servers and banners to use for the authentication proxy and to configure non-default timeout values. The authentication proxy for IOS devices is a service that forces users to log in and authenticate when trying to make HTTP, Telnet, or FTP connections through an IOS device. The settings you configure here work in conjunction with your AAA rules; only if a AAA rule requires user authentication for one of these services does your AuthProxy settings come into play.

Ensure that your configuration of this policy is consistent with your **Firewall > AAA Rules** policy. Additionally, you must use the **Platform > Device Admin > AAA** policy to define the AAA server groups to use for authenticating user access; this policy defines only the authorization and accounting server groups. If you also want to use authorization proxy for HTTPS access, you must enable SSL and configure AAA in the **Platform > Device Admin > Device Access > HTTP** policy in addition to enabling HTTP authorization proxy in your AAA rules policy.

Tip You must configure per-user ACLs in your AAA server to define the privileges you want to apply to each user. When configuring authorization, specify **AAA** as the service (e.g. service = AAA), with a privilege level of 15. For more information on configuring the AAA server, including information on configuring authentication proxy in general, see the "Configuring the Authentication Proxy" section in the *Cisco IOS Security Configuration Guide: Securing User Services, Release 12.4T* at http://www.cisco.com/en/US/docs/ios/sec_user_services/configuration/guide/sec_cfg_authen_prxy_ps6441_TSD_Products_Configuration_Guide_Chapter.html.

Navigation Path

To access the AAA page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > AAA** from the Policy selector.
- (Policy view) Select **Firewall > Settings > AAA** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Settings > AAA**.

Related Topics

- Understanding AAA Rules, page 15-1
- Understanding How Users Authenticate, page 15-2
- Configuring AAA Rules for IOS Devices, page 15-7

Field Reference

Table 15-8 AAA Firewall Settings Policy

Element	Description
Virtual IP Address	You use the Virtual IP Address only in communications between the IOS HTTP authentication and clients. For the system to operate correctly, the virtual IP address must be set (it cannot be 0.0.0.0), and no other device on the network can have the same address. Configure with an unassigned and unused gateway IP address, such as 1.1.1.1.

Table 15-8 AAA Firewall Settings Policy (continued)

Element	Description
General Tab	
Authorization Server Groups	The AAA server group policy objects that identify the LDAP, TACACS+, or RADIUS servers that will provide per-user authorization control. You can also use the LOCAL user database defined on the device. Enter the names of the server group objects, or click **Select** to select them from a list or to create new objects. Ensure that you put the groups in priority order; authorization is attempted with the first group and if that group is not available, with subsequent groups.
Accounting Server Groups Use Broadcast for Accounting	The AAA server group policy objects that identify the LDAP, TACACS+, or RADIUS servers that will provide accounting services. Accounting collects per-user usage information for billing, security, or resource allocation purposes. Enter the names of the server group objects, or click **Select** to select them from a list or to create new objects. Ensure that you put the groups in priority order; if you do not select the broadcast option, accounting is attempted with the first group and if that group is not available, with subsequent groups. If you select **Use Broadcast for Accounting**, accounting records are sent simultaneously to the first server in each group. If the first server is unavailable, failover occurs using the backup servers defined within that group.
Accounting Notice	The types of accounting notices to be sent to the accounting server groups: • Start-stop—Sends a start accounting notice at the beginning of a user process and a stop accounting notice at the end of the process. The start accounting record is sent in the background. The requested user process begins regardless of whether the start accounting notice was received by the accounting server. • Stop-only—Sends a stop accounting notice at the end of the requested user process. • None—No accounting records are sent.
HTTP Banner FTP Banner Telnet Banner	The banner you want to present on the authentication proxy page to the user when the user is prompted to authenticate for the specified service: • Disable Banner Text—No banner is displayed. • Use Default Banner Text—Displays the default banner “Cisco Systems, *router hostname* Authentication.” • Use Custom Banner Text—Enter the text you want to present to the user.

Table 15-8 AAA Firewall Settings Policy (continued)

Element	Description
Use HTTP banner from File URL	Whether you want to use your own web page to authenticate HTTP connections. Enter the URL for your HTTP banner. If you configure both HTTP banner text and a URL, the URL banner take precedence; however, the banner text is also configured on the device.
Advanced Tab	
Global Inactivity Time	The length of time, in minutes, that the authentication proxy for a user is maintained when there is no user activity in the session. If this timer expires, the user session is cleared along with its dynamic user access control list (ACL), and the user must re-authenticate. The range is 1 to 2,147,483,647. The default is 60 minutes. Ensure that this timeout value is greater than or equal to the idle timeout values configured in the **Firewall > Settings > Inspection** policy; otherwise, timed-out user sessions might continue to be monitored and eventually hang.
Global Absolute Time	The length of time, in minutes, that an authentication proxy user session can remain active. After this timer expires, the user session must go through the entire process of establishing its connection as if it were a new request. The range is 0 to 35,791. The default is 0, which means that there is no global absolute timeout; user sessions are maintained as long as they are active.
Interface Timeout Table	This table contains the interfaces for which you want to configure timeout values that differ from the global timeout values. If you want to use the global values for all interfaces, you do not need to configure anything in this table. • To add an interface with customized timeout values, click the **Add Row** button and fill in the Firewall AAA IOS Timeout Value Setting, page 15-30. • To edit a setting, select it and click the **Edit Row** button. • To delete a setting, select it and click the **Delete Row** button.

Firewall AAA IOS Timeout Value Setting

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the Firewall AAA IOS Timeout Value Setting dialog box to configure idle and absolute timeout values for specific interfaces. These values override the global timeout values configured on the **Firewall > Settings > ScanSafe Web Security** policy Server Timeout tab.

Navigation Path

From the Advanced tab of the AAA Page, page 15-28, click the **Add Row** button beneath the table of interfaces, or select a row and click the **Edit Row** button.

Field Reference

Table 15-9 Firewall AAA IOS Timeout Value Setting Dialog Box

Element	Description
Interfaces	The interfaces or interface roles for which you are configuring timeout values. Enter the names of the interfaces or roles, or click **Select** to select them from a list or to create new interface roles. Separate multiple entries with commas.
Auth Proxy Tab	
Inactivity/Cache Time	The length of time, in minutes, that the authentication proxy for a user is maintained when there is no user activity in the session on the interface. If this timer expires, the user session is cleared along with its dynamic user access control list (ACL), and the user must re-authenticate. The range is 1 to 2,147,483,647. The default is the global inactivity timeout value (whose default is 60 minutes).
Absolute Time	The length of time, in minutes, that an authentication proxy user session can remain active on the interface. After this timer expires, the user session must go through the entire process of establishing its connection as if it were a new request. The range is 1 to 35,791. The default is 0, which means that there is no absolute timeout; user sessions are maintained as long as they are active.
Authentication Proxy Method (IOS)	The protocols to which these timeout values should apply. You can select any combination of HTTP, FTP, or Telnet.

Table 15-9 Firewall AAA IOS Timeout Value Setting Dialog Box (continued)

Element	Description
HTTP/NTLM Tab	The HTTP and the NTLM areas contain the same following fields and selections: Set the **Inactivity/Cache Time** and the **Absolute Time** for HTTP/NTLM and then, if desired, select Enable Passive Authentication. Finally, select the **Identity Policy** that you want to apply.
Method Order Tab	Select the checkbox of each method you want to employ and then use the up and down arrows to arrange the methods in the order you desire.
AAA Settings Tab	Select the **AAA Settings** tab to specify the Authentication, Authorization, and Account settings as detailed below.
Authenticate Using	In the Authenticate Using section you can select the server group(s) to use for authentication. Choices are: • **None**—No authentication • **Default**—Use the default authentication server group(s). • **Custom**—Enable the selection of user specified authentication server group(s). Then click **Select** to specify or add a server group.
Authorize Exec Operation Using	In the Authorize Exec Operation Using section you can select the server group(s) to use for authorization of executive operations. Choices are: • **None**—No authorization • **Default**—Use the default authorization server group(s) • **Custom**—Enable the selection of user specified authorization server group(s). Then click **Select** to specify or add a server group.
Perform Exec Operation Using	In the Authorize Exec Operation Using section you can select the server group(s) to use for performing executive operations. Choices are: • **None**—No authorization • **Default**—Use the default server group(s) • **Custom**—Enable the selection of user specified server group(s). Then click **Select** to specify or add a server group.
Accounting Notice	Use Accounting Notice to specify accounting operations. • **None**—No accounting notices • **Start-stop**—Accounting notices at the beginning and end of operations • **Stop-only**—Accounting notices at the end of operations.
Accounting Server Groups	Specify what accounting server groups to use. Either enter or select the Accounting Server group. **Note** If you choose to select an accounting server group, you are also give the option to add an Accounting Server group.
Use Broadcast for Accounting	Select this checkbox to broadcast accounting notices.

Managing Firewall Access Rules

Access rules define the rules that traffic must meet to pass through an interface. When you define rules for incoming traffic, they are applied to the traffic before any other policies are applied (with the exception of less common AAA rules). In that sense, they are your first line of defense.

Tip For some types of devices, you can configure IPv6 access rules in addition to IPv4 access rules. For information on supported device types, see IPv6 Support in Security Manager, page 1-8.

The following topics help you understand and work with access rules:

- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7
- Configuring Expiration Dates for Access Rules, page 16-22
- Configuring Settings for Access Control, page 16-23
- Using Automatic Conflict Detection, page 16-28
- Viewing Hit Count Details, page 16-36
- Importing Rules, page 16-40
- Optimizing Access Rules Automatically During Deployment, page 16-46

The following topics can help you with general rule table usage:

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19

Understanding Access Rules

Access rules policies define the rules that allow or deny traffic to transit an interface. Typically, you create access rules for traffic entering an interface, because if you are going to deny specific types of packets, it is better to do it before the device spends a lot of time processing them.

When you deploy access rules to devices, they become one or more entries (ACEs) to access control lists (ACLs) that are attached to interfaces. Typically, these rules are the first security policy applied to packets; they are your first line of defense. You use access rules to filter out undesired traffic based on service (protocol and port numbers) and source and destination addresses, either permitting the traffic or denying (dropping) it. Each packet that arrives at an interface is examined to determine whether to forward or drop the packet based on criteria you specify. If you define access rules in the out direction, packets are also analyzed before they are allowed to leave an interface.

Tip For ASA 8.3+ devices, you can augment interface-specific access rules with global access rules. For more information, see Understanding Global Access Rules, page 16-3.

When you permit traffic in an access rule, subsequent policies might end up dropping it. For example, inspection rules, web filter rules, and zone-based firewall rules are applied after a packet makes it through the interface's access rules. These subsequent rules might then drop the traffic based on a deeper analysis of the traffic; for example, the packet header might not meet your inspection requirements, or the URL for a web request might be for an undesired web site.

Thus, you should carefully consider the other types of firewall rules you intend to create when you define access rules. Do not create a blanket denial in an access rule for traffic that you really want to inspect. On the other hand, if you know that you will never allow a service from or to a specific host or network, use an access rule to deny the traffic.

Keep in mind that access rules are ordered. That is, when the device compares a packet against the rules, it searches from top to bottom and applies the policy for the first rule that matches it, and ignores all subsequent rules (even if a later rule is a better match). Thus, you should place specific rules above more general rules to ensure those rules are not ignored. To help you identify cases where IPv4 rules will never be matched, and to identify redundant rules, you can use the automatic conflict detection and policy query tools. For more information, see Using Automatic Conflict Detection, page 16-28 and Generating Policy Query Reports, page 12-28.

The following are additional ways in which you can evaluate your access rules:

- Combine rules—You can use a tool to evaluate your IPv4 rules and combine them into a smaller number of rules that perform the same functions. This can leave you with a smaller, easier to manage list of rules. For more information, see Combining Rules, page 12-22.
- Generate hit counts—You can use a tool to view the hit count statistics maintained by the device for IPv4 and IPv6 ACLs. This can tell you how often a rule has permitted or denied traffic. For more information, see Viewing Hit Count Details, page 16-36.
- View events collected by CS-MARS—You can analyze real time or historical events related to an IPv4 rule using the Cisco Security Monitoring, Analysis and Response System application if you configured it to monitor the device and you configure the rule to generate syslog messages. For more information, see Viewing CS-MARS Events for an Access Rule, page 72-42.

For more conceptual information on access rules, see the following topics:

- Understanding Global Access Rules, page 16-3
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5

Related Topics

- Configuring Access Rules, page 16-7
- Configuring Expiration Dates for Access Rules, page 16-22

- Configuring Settings for Access Control, page 16-23
- Expanding Object Groups During Discovery, page 12-35
- Importing Rules, page 16-40
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19

Understanding Global Access Rules

Traditionally, access rules (ACLs), which control which traffic can flow through a device, are applied to device interfaces. However, with ASA devices running software release 8.3+, you have the option to create global access rules for IPv4 and IPv6.

Global access rules are defined as a special ACL that is processed for every interface on the device for traffic entering the interface. Thus, although the ACL is configured once on the device, it acts like a secondary interface-specific ACL defined for the In direction. (Global rules are always for the In direction, never the Out direction.)

When traffic enters an interface on an ASA 8.3+ device, when applying ACLs, the device first applies any interface-specific access rules to the traffic. It then applies global rules. (Overall processing is explained in Understanding the Processing Order of Firewall Rules, page 12-2.)

Global rules are best used for rules that you want to apply to all traffic that enters a device regardless of which interface it enters. For example, there might be a specific host or subnet that you always want to deny or permit. You can create these as global rules, so they are configured once on the device instead of configured again for each interface (although functionally the same as an interface-specific rule configured for the All-Interfaces role, All-Interfaces rules are repeated for each interface rather than being configured once on the device).

Tip If you want to configure the same set of global rules for more than one device, create a shared policy and inherit it in the IPv4 or IPv6 access rules policy for each device. Ensure that all global rules are in the Default section of the shared policy. If you put any global rules in the Mandatory section, you will not be able to inherit it on devices that have local interface-specific access rules defined. For more information on shared and inherited policies, see Local Policies vs. Shared Policies, page 5-3 and Understanding Rule Inheritance, page 5-4.

When you configure access rules for an ASA 8.3+ device in Security Manager, interface-specific and global rules are configured in the same policy. However, because the device always processes interface-specific rules first, Security Manager prevents you from intermixing these different types of rules. Therefore, if you configure both interface-specific and global rules on a device, keep the following in mind:

- Global rules always come last in the access rules policy. All interface-specific rules come before global rules.
- You cannot move rules in a way that violates the required order. For example, you cannot move an interface-specific rule below a global rule, or a global rule above an interface-specific rule.
- You cannot create rules in a location that violates the required order. For example, if you select an interface-specific rule, and another interface-specific rule follows it in the table, you cannot create a global rule. If you try to create the wrong kind of rule, when you save the rule, Security Manager

will ask you if the rule can be created at the nearest valid location. You must accept the suggestion or the rule will not be added to the table. You can always move the rule after creating it if the suggested location is not ideal (but without violating the rules on order).

- You cannot inherit a policy if the rules in the inherited policy will violate the required order. For example, if you create global rules in the device policy, and try to inherit a shared policy that contains interface-specific rules in the Default section, Security Manager will prevent you from inheriting the policy.
- After assigning or inheriting a shared policy, you cannot edit the policy in a way that will violate rule order on any device that uses the policy.
- If you assign or inherit a policy that contains global rules on a device that does not support them, all global rules are ignored and not configured on the device. For example, if you permit all traffic from host 10.100.10.10 in a global rule in a shared policy, and assign that policy to an IOS device, the rule permitting 10.100.10.10 access is not configured on the IOS device, and traffic from that host is handled either by another interface-specific policy, or the default deny all policy. As a good practice, you should not assign shared policies that contain global rules to devices that do not support them, to ensure that you do not mistakenly believe the policy defined in a global rule is being configured on the unsupported device.

There are also some changes in how certain tools work with global rules:

- Find/Replace—You can search for global rules by using the Global interface name. However, there is no way to convert between global and interface-specific rules. Although you can find global rules using the Global interface name, if you try to replace an interface name with the name "Global," you are actually creating an interface-specific access rule that uses a policy object named Global.
- Rule Combiner—Interface-specific and global rules are never combined.

Related Topics

- Understanding Access Rules, page 16-1
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7
- Moving Rules and the Importance of Rule Order, page 12-19

Understanding Device Specific Access Rule Behavior

If you do not create an access rule policy, the following is the default behavior based on the type of device, and what happens when you create an access rule:

- IOS devices—Permit all traffic through an interface.

  When you create an access rule permitting source A to destination B without configuring TCP/UDP inspection on the inspection rule table, or configuring the **established** advanced option on the rule, the device permits any packet from A to B. However, for any returning packet from B to A, the packet is not allowed, unless there is a corresponding access rule permitting that packet. If you configure TCP/UDP inspection on the traffic the inspection rule table, a rule permitting B to A is not needed in the access rule, as any returning packet from B to A automatically passes the device.
- ASA and PIX devices—Permit traffic from a higher-security interface to a lower-security interface. Otherwise, all traffic is denied.

If an access rule allows TCP/UDP traffic in one direction, the appliance automatically allows return traffic (you do not need to configure a corresponding rule for the return traffic), except for ICMP traffic, which does require a return rule (where you permit the reverse source and destination), or you must create an inspection rule for ICMP.

- FWSM devices—Deny all traffic entering an interface, permit all traffic leaving an interface.

  You must configure access rules to allow any traffic to enter the device.

If you create any rules for an interface for any type of device, the device adds an implicit **deny any** rule at the end of the policy. It is a good practice for you to add this rule yourself so that you remember it is there. Adding the rule also allows you to get hit count information for the rule. For more information, see Viewing Hit Count Details, page 16-36.

Tip When you create the access rule policy, ensure that you include a rule that will permit access to the device from the Security Manager server, or you will not be able to manage the device using the product.

Related Topics

- Understanding Access Rules, page 16-1
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5

Understanding Access Rule Address Requirements and How Rules Are Deployed

One of the complexities of creating access control lists using the operating system commands on the command line interface (CLI) is the fact that different operating systems have different IP address formats for source and destination addresses.

For example, Cisco IOS Software requires that you enter addresses using wildcard masks instead of subnet masks. To create a rule for the 10.100.10.0/24 network (subnet mask 255.255.255.0), you are required to enter the address as 10.100.10.0 0.0.0.255. The 0 and 1 have the reverse meaning in a wildcard mask that they have in a subnet mask. In ASA, PIX, and FWSM software, however, you use subnet masks, so you enter 10.100.10.0 255.255.255.0.

Security Manager simplifies addressing requirements for access rules: you always use the subnet mask. You are not even allowed to enter a wildcard mask. When you deploy the access rules to a device, Security Manager considers the operating system of the device and converts subnet masks to wildcard masks automatically when needed.

This makes it possible for you to create shared rules based on logical policies and to apply those rules to all of your devices. For example, there might be a set of access rules that you want all devices to use, in which case you can create the shared policy and assign it as the inherited policy for all devices. You do not have to worry about defining rules using the “right” syntax for the device type. You can use the same network/host objects that you use in other types of policies to identify targeted hosts and networks.

The specific CLI commands generated in deployed configurations are also based on the type of device. For IOS devices, the **ip access-list** command is used. For ASA, PIX, FWSM devices, the **access-list** or **ipv6 access-list** command is used and the **access-group** command binds it to the interface. With ASA, PIX, FWSM, and IOS 12.4(20)T+ devices, if you use network/host objects to identify the source or destination addresses for a rule, the **object-group** command is used to create object groups for those network/host objects. Object groups are also created for service objects.

Tips

- Because you can use network/host objects to identify a source or destination, and you can configure deployment optimization for rules, there is not always a one-to-one relationship between an access rule and ACEs in the CLI definition of an ACL.
- All access lists created from firewall rules are extended access lists (rather than standard). Security Manager applies a system-generated name to the ACL unless you specify a name for the ACL on the Access Control Settings Page, page 16-24. The name applies to the ACL that includes all of the rules related to the interface and direction for which the name is defined.
- There are several deployment options that control how object groups are deployed. This topic describes the default behavior. On the Deployment Page, page 11-13 (select **Tools > Security Manager Administration > Deployment**), you can deselect the option to create object groups from network/host objects. You can also optimize object groups during deployment (see Optimizing Network Object Groups When Deploying Firewall Rules, page 12-35), create new object groups from rules with multiple services or source and destination addresses, or remove unused object groups.
- The deployment options also include settings that control the names of ACLs generated from access rules and how many ACLs are created. By default, Security Manager creates a unique ACL for each interface, even if this means that several duplicate ACLs are created.

  If you select **Enable ACL Sharing for Firewall Rules**, Security Manager can create a single ACL and apply it to multiple interfaces, thus avoiding the creation of unnecessary duplicate ACLs. However, ACL sharing occurs only if it can be done while preserving your ACL naming requirements:

  - If you specify an ACL name for an interface and direction, that name is always used, even if it means a duplicate ACL must be created. For more information, see Configuring Settings for Access Control, page 16-23.
  - If you select **Reuse existing names** for the Firewall Access-List Names property, the existing names are preserved (unless you override them in the access control settings policy). This means that you might end up with duplicate ACLs under different names if duplicate ACLs already exist on the device.

  Tip: To maximize ACL sharing, ensure that you select **Reset to CS-Manager Generated Names** for the Firewall Access-List Names property, select **Speed** for the Optimize the Deployment of Access Rules For property, and that you do not configure ACL names in the access control settings policy.

  For more detailed information about the **Enable ACL Sharing for Firewall Rules** property, see Deployment Page, page 11-13.
- IPv4 and IPv6 ACLs cannot have the same name.

Related Topics

- Understanding Access Rules, page 16-1
- Configuring Access Rules, page 16-7
- Configuring Settings for Access Control, page 16-23
- Expanding Object Groups During Discovery, page 12-35

Configuring Access Rules

Access rules policies define the rules for allowing traffic to pass through an interface. If you do not configure an access rules policy, the device behavior differs based on device type as explained in Understanding Device Specific Access Rule Behavior, page 16-4.

Note Prior to the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, separate pages, policies and policy objects were provided for configuring IPv4 and IPv6 firewall rules and policies. With Security Manager 4.4 and ASA 9.0+, these policies and policy objects were "unified," meaning one set of rules in which you can use either IPv4 or IPv6 addresses, or a mixture of both. (See Policy Object Changes in Security Manager 4.4, page 1-10 for additional information.) However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided in Device view, while in Policy view, IPv4 and unified versions of the access-rule policy types are provided. In addition, a utility that you can use to convert existing IPv4 policies is provided (see Converting IPv4 Rules to Unified Rules, page 12-28). The following descriptions apply to all versions of the access rule table, except where noted.

If you assign an IPv4 access-rule shared policy to a 9.0+ device, you will no longer be able to assign unified versions of those policies to that device. Likewise, if you assign a unified access-rule shared policy to a 9.0+ device, you will no longer be able to assign IPv4 versions of those shared policies to that device--the device will not be included in the list of available devices on the Assignments tab for the shared policy.

Before you configure access rules, consider the other types of firewall rules you will configure. Access rules are processed before all other types of rules except AAA rules. See the following topics for more information about things you should consider:

- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5

Before You Begin

You might have a set of access rules that you want to apply to all devices. To do this, you can create a shared rule and inherit its rules to each device's access rules policy. For more information, see Creating a New Shared Policy, page 5-54 and Inheriting or Uninheriting Rules, page 5-47.

Related Topics

- Using Sections to Organize Rules Tables, page 12-20
- Copying Policies Between Devices, page 5-33
- Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37
- Understanding Networks/Hosts Objects, page 6-80
- Understanding Interface Role Objects, page 6-73
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

Step 1 Do one of the following to open the Access Rules Page, page 16-10:

- (Device view) Select **Firewall > Access Rules** (or **Firewall > Settings > IPv6 Access Rules**) from the Policies selector.

- (Policy view) Select **Firewall > Access Rules** (or **Firewall > Settings > IPv6 Access Rules**) from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add and Edit Access Rule Dialog Boxes, page 16-14.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10. Special rules apply if you mix interface-specific and global rules in a policy; for more information, see Understanding Global Access Rules, page 16-3.

Step 3 Configure the rule. Following are the highlights of what you typically need to decide while configuring your rule. For specific information on configuring the fields, see Add and Edit Access Rule Dialog Boxes, page 16-14.

- Permit or Deny—Whether you are allowing traffic that matches the rule or dropping it.
- Source and Destination addresses—If the rule should apply no matter which addresses generated the traffic or their destinations, use "All-Addresses" as the source or destination. If the rule is specific to a host or network, enter the addresses or network/host objects. For information on the accepted address formats, see Specifying IP Addresses During Policy Definition, page 6-87.
- Source and Destination Security Groups (ASA 9.0+ only)—You can specify TrustSec security groups used to filter traffic in addition to the source and destination addresses. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups.
- Source Users (ASA 8.4.2+ only)—You can further define the traffic source by specifying Active Directory (AD) user names (in the format NetBIOS_DOMAIN\username), user groups (NetBIOS_DOMAIN\\user_group), or identity user group objects that define the names and groups. The user specification is conjoined to the source address to limit the match to user addresses within the source address range. For more information, see Configuring Identity-Based Firewall Rules, page 13-21 and Creating Identity User Group Objects, page 13-19.
- Services—Use the IP service to apply to any traffic (for example, if you want to deny all traffic from a specific source). Otherwise, select the more specific service, which is a protocol and port combination, that you are targeting.
- Interfaces or Global—The interface or interface role for which you are configuring the rule, or Global to create global access rules on ASA 8.3+ devices (see Understanding Global Access Rules, page 16-3).
- Advanced settings—Click **Advanced** to open the Advanced dialog box for configuring additional settings. You can configure the following options; for detailed information, see Advanced and Edit Options Dialog Boxes, page 16-17.
  - Logging options. If you are using Security Manager or CS-MARS to monitor the device, ensure that you enable logging.
  - The direction of traffic to which this rule should apply (in or out). The default is in. You cannot change this setting for global rules.
  - The time range for the rule, which allows you to configure rules that work only for specific periods of time, such as during working hours. For more information, see Configuring Time Range Objects, page 6-71.
  - IOS device options for fragmentation and allowing the return of traffic for established outbound sessions.

– Rule expiration dates and notification settings. For more information, see Configuring Expiration Dates for Access Rules, page 16-22.

Step 4 Click **OK** when you are finished defining your rule.

Note You can enable conflict detection (see Enable Conflict Detection) to see if the new rule conflicts or overlaps with other rules. For more information, see Using Automatic Conflict Detection, page 16-28.

Note While adding or editing a rule, any two rules might become identical (example shown as 1 and 2 in Figure 16-1) except for differences in time range and/or logging values (defined in Advanced and Edit Options Dialog Boxes)—

- Cisco Security Manager deploys only the rule that is at the bottom (2 in Figure 16-1).
- Only rule (2) is used to identify configuration changes in preview configuration (see Previewing Configurations, page 8-44).
- If rule (2) has been deployed on the device, the preview configuration will not detect any changes.

Figure 16-1 Identical Rules

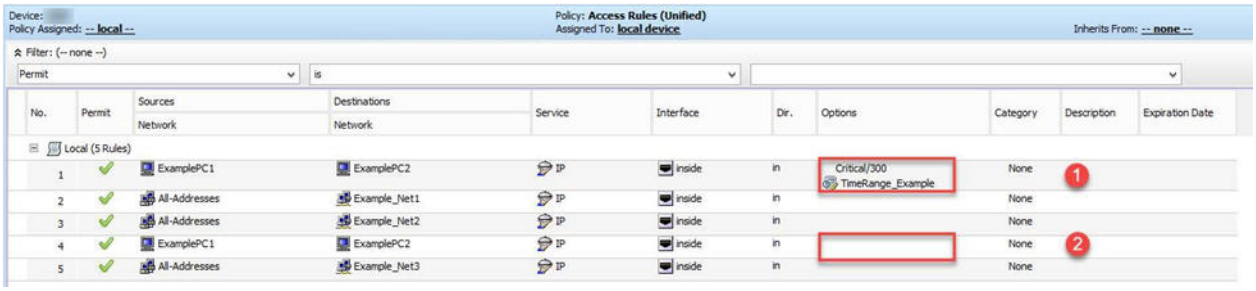

Step 5 If you are required to find the bottom rule (example, (2) in Figure 16-1), that prevents Cisco Security Manager from deployment of a top rule (1 in Figure 16-1), do the following:

a. Enable Conflict Detection for the device.

b. Generate Report for the conflicts found.

c. In the report, under the Rule No column, find the bottom rule (2), determine the rule number it conflicts with [rule (1)] and delete rule (1), if required.

Step 6 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. For more information, see Moving Rules and the Importance of Rule Order, page 12-19. There are special restrictions for moving rules when you mix interface-specific and global rules; see Understanding Global Access Rules, page 16-3.

Step 7 If you already have a large number of rules, consider analyzing and combining them before deploying the new rules. You can use the conflict detection tool to analyze your rules (see Using Automatic Conflict Detection, page 16-28). If analysis shows that you have a lot of redundant rules, right-click anywhere in

the rules table and choose **Combine Rules** to combine them. You can either allow Security Manager to evaluate all rules for combination, or just the rules you select before starting the rule combination tool. For more information, see Combining Rules, page 12-22.

Access Rules Page

Use the Access Rules page to configure access control rules for device interfaces. Access rules policies define the rules that allow or deny traffic to transit an interface. Typically, you create access rules for traffic entering an interface, because if you are going to deny specific types of packets, it is better to do it before the device spends a lot of time processing them. Access rules are processed before other types of firewall rules.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate policies and objects for configuring IPv4 and IPv6 access rules were "unified," meaning one set of access rules in which you can use either IPv4 or IPv6 addresses, or a mixture of both. (See Policy Object Changes in Security Manager 4.4, page 1-10 for additional information.) In Policy view, IPv4 and unified versions of the access policy type are provided. In addition, a utility that you can use to convert existing IPv4 policies is provided (see Converting IPv4 Rules to Unified Rules, page 12-28). The following descriptions apply to apply to all versions of the access rule table, except where noted.

Read the following topics before you configure access rules:

- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7

Tip Disabled rules are grayed out. When you deploy the configuration, disabled rules are removed from the device. For more information, see Enabling and Disabling Rules, page 12-20.

Navigation Path

To open the Access Rules page, do one of the following:

- (Device view) Select a device, then select **Firewall > Access Rules** (or **Firewall > Settings > IPv6 Access Rules**) from the Policies selector.
- (Policy view) Select **Firewall > Access Rules** (or **Firewall > Settings > IPv6 Access Rules**) from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Policies > Access Rules** (or **Edit Firewall Policies > IPv6 Access Rules**).

Related Topics

- Configuring Expiration Dates for Access Rules, page 16-22
- Configuring Settings for Access Control, page 16-23
- Adding and Removing Rules, page 12-9

- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19
- Using Sections to Organize Rules Tables, page 12-20
- Using Rules Tables, page 12-8
- Filtering Tables, page 1-48

Field Reference

Note For details on the fields and user interface elements available as part of the automatic conflict detection feature, see Understanding the Automatic Conflict Detection User Interface, page 16-30.

Table 16-1 Access Rules Page

Element	Description
Expand all rows/Collapse all rows	Use these buttons to expand or collapse all sections in the rules table. **Note** The buttons are located in the upper-right corner of the Filter area above the access rules table.
Conflict Indicator icons	Identifies conflicts and provides a quick visual representation of the type of conflict. For more details, including types of conflicts and the actions you can take from this column, see Understanding the Automatic Conflict Detection User Interface, page 16-30.
No.	The ordered rule number.
Permit	Whether a rule permits or denies traffic based on the conditions set: • Permit—Shown as a green check mark. • Deny—Shown as a red circle with slash.
Sources	The sources of traffic for this rule; can be networks, security groups (ASA 9.0+ only), and users. Multiple entries are displayed on separate lines within the table cell.
Destinations	The destinations for this rule; can be networks and security groups (ASA 9.0+ only). Multiple entries are displayed on separate lines within the table cell.
Service	The services or service objects that specify the protocol and port of the traffic to which the rule applies. Multiple entries are displayed on separate lines within the table cell. See Understanding and Specifying Services and Service and Port List Objects, page 6-100.

Table 16-1 Access Rules Page (continued)

Element	Description
Hit Count	Number of times this rule has been "hit"; that is, number of times it has permitted or denied traffic; it is actually the sum of the hit counts for all accesss control entries (ACEs) created by the rule. This information is useful in debugging the deployed policies. Use the Refresh Hit Count button at the bottom of this page to update the hit information; opens the Hit Count Selection Summary Dialog Box, page 16-20. **Note** The Hit Count of a duplicated ACE, either within the same rule or different rules, is always set to 0. You can right-click this cell and choose Show Hit Count Details to open the Hit Count Details pane at the bottom of the Configuration Manager window. See Viewing Hit Count Details, page 16-36 for more information.
Last Hit Time	Timestamp for the most-recent hit.
Interface	The interfaces or interface roles to which the rule is assigned. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. Multiple entries are displayed on separate lines within the table cell. See Understanding Interface Role Objects, page 6-73. For ASA 8.3+ devices, global rules are indicated with the name **Global** and a special icon to distinguish them from rules that use interface or interface role names (for an explanation of the icons, see Specifying Interfaces During Policy Definition, page 6-76).
Dir.	The direction of the traffic to which this rule applies: • In—Packets entering the interface. • Out—Packets exiting the interface.
Options	Any additional options configured for the rule. These include logging, time range, and some additional IOS rule options. See Advanced and Edit Options Dialog Boxes, page 16-17.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of the rule, if any.
Expiration Date	The date on which the rule expires. Expired rules show **Expired** in bold text. Expired rules are not automatically deleted.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Page elements below the rules table	

Table 16-1 Access Rules Page (continued)

Element	Description
Enable conflict detection Generate Report (neither option presented on the IPv6 Access Control page)	Enable or disable automatic conflict detection. This feature is enabled by default and the setting is managed per user. Disabling conflict detection for one access rules table will also disable the feature for other access rules tables. You can disable conflict detection while creating your rule table or making large modifications and then re-enable it when you are ready to verify your changes. See Using Automatic Conflict Detection, page 16-28. **Note** For details on the fields and user interface elements available as part of the automatic conflict detection feature, see Understanding the Automatic Conflict Detection User Interface, page 16-30. If conflict detection is enabled, you can click the **Generate Report** button to create an HTML report of any rule conflicts detected. This report can be printed or exported to another application.
Refresh Hit Count (not presented on the IPv6 Access Control page)	Click this button to update the hit information displayed in the table; opens the Hit Count Selection Summary Dialog Box, page 16-20.
Query (not presented on the IPv6 Access Control page)	Click this button to run a policy query, which can help you evaluate your rules and identify ineffective rules. See Generating Policy Query Reports, page 12-28
Find and Replace button (binoculars icon)	Click this button to search for various types of items within the table and to optionally replace them. See Finding and Replacing Items in Rules Tables, page 12-16.
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope or section. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Add Row button	Click this button to add a rule to the table after the selected row using the Add and Edit Access Rule Dialog Boxes, page 16-14. If you do not select a row, the rule is added at the end of the local scope. For more information about adding rules, see Adding and Removing Rules, page 12-9.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells. For more information, see Editing Rules, page 12-10.
Delete Row button	Click this button to delete the selected rule.

Right-click Menu

A right-click menu is also available. This menu provides access to many of the functions listed above; the options presented depend on the location right-clicked:

- If you right-click a rule in the table, the options may include editing functions relative to the specific table cell right-clicked. For example, the command “Show Hit Count Details” is included when you right-click a Hit Count cell. See Editing Rules, page 12-10 for more information.

- You can also navigate from a rule to events associated with that rule in either Event Viewer or CS MARS. For more information, see Viewing Events for an Access Rule, page 69-56 and Viewing CS-MARS Events for an Access Rule, page 72-42.
- The Import Rules and Combine Rules options are also included in the right-click menu. See Importing Rules, page 16-40 and Combining Rules, page 12-22 for more information about these options.

Add and Edit Access Rule Dialog Boxes

Use the Add and Edit Access Rule dialog boxes to add and edit security-device access rules.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate pages for configuring IPv4 and IPv6 access rules were unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided. The following descriptions apply to apply to all versions of the page, except where noted.

Read the following topics before you configure access rules:

- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7

Navigation Path

On the Access Rules Page, page 16-10, click the **Add Row** button or select a row and click the **Edit Row** button.

Note Prior to Cisco Security Manager 4.13, the **Add Access Rule** dialog was populated with default values. Starting from 4.13, the user can customize the appearance of default values by updating the **csm.properties** file. For more information, see Customizing defaults in the Add Access Rule dialog, page 16-48

Related Topics

- Configuring Expiration Dates for Access Rules, page 16-22
- Editing Rules, page 12-10
- Adding and Removing Rules, page 12-9
- Importing Rules, page 16-40
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

Field Reference

Table 16-2 Add and Edit Access Rule Dialog Boxes

Element	Description
Enable Rule	Check this box to enable this rule, which means the rule becomes active when you deploy the configuration to the device. Deselect to disable the rule while keeping the rule definition. Disabled rules are shown overlaid with hash marks in the rule table. See Enabling and Disabling Rules, page 12-20 for more information.
Action	Whether the rule permits or denies traffic based on the conditions you define.

Table 16-2 Add and Edit Access Rule Dialog Boxes (continued)

Element	Description
Sources	Provide traffic sources for this rule; can be networks, security groups, and users. You can enter values or object names, or Select objects, for one or more of the following types of sources: • Network – You can specify a various network, host and interface definitions, either individually or as objects. If you Select an interface object as a source, the dialog box displays tabs to differentiate between hosts/networks and interfaces. The "All-Address" objects do not restrict the rule to specific hosts, networks, or interfaces. These addresses are IPv4 or IPv6 addresses for hosts or networks, network/host objects, interfaces, or interface roles. **Note** You can only specify a fully qualified domain name (FQDN) by providing an FQDN network/host object, or a group object that includes an FQDN object. You cannot directly type in an FQDN. See Understanding Networks/Hosts Objects, page 6-80, Specifying IP Addresses During Policy Definition, page 6-87 and Understanding Interface Role Objects, page 6-73 for additional information about these definitions. • Security Groups (ASA 9.0+) – Enter or Select the name or tag number for one or more source security groups for the rule, if any. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. • Users – Enter or Select the Active Directory (AD) user names, user groups, or identity user group objects for the rule, if any. You can enter any combination of the following: – Individual user names: NetBIOS_DOMAIN\username – User groups (note the double \): NetBIOS_DOMAIN\\user_group – Identity user group object names. For more information, see: – Selecting Identity Users in Policies, page 13-21 – Configuring Identity-Based Firewall Rules, page 13-21 – Creating Identity User Group Objects, page 13-19 **Note** Enter more than one value in any of these fields by separating the items with commas. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.

Table 16-2 Add and Edit Access Rule Dialog Boxes (continued)

Element	Description
Destinations	Provide traffic destinations for this rule; can be networks or security groups. As with Sources, you can enter values or object names, or Select objects, for one or more destinations of Network and Security Group (ASA 9.0+) type.
Services	The services that define the type of traffic upon which to act. You can enter or Select any combination of service objects and service types (which are typically a protocol and port combination). Enter more than one value by separating the items with commas. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.
Interfaces Global (ASA 8.3+)	Specify whether you are creating an interface-specific or global rule. Global rules are available only for ASA 8.3+ devices, and are handled according to special rules (for detailed information, see Understanding Global Access Rules, page 16-3). If you select **Interfaces**, enter or Select the name of the interface or the interface role to which the rule is assigned, or click **Select** to select the interface or role from a list. An interface must already be defined to appear on the list. For bridge groups in routed mode, you can create access rules for both the Bridge Virtual Interface (BVI) and each bridge group member interface. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. See Understanding Interface Role Objects, page 6-73. Global rules are created as a special global ACL that is not attached to specific interfaces, but are processed for all interfaces in the In direction after interface-specific rules.
Description	An optional description of the rule (up to 1024 characters).
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Advanced button	Click this button to configure other settings for the rule, including logging configuration, traffic direction, time ranges, and rule expiration dates. For more information, see Advanced and Edit Options Dialog Boxes, page 16-17.

Advanced and Edit Options Dialog Boxes

Use the Advanced dialog box to configure additional settings for an access rule. These settings are displayed in three different cells of the access-rule table: direction, options, and rule expiration. You can then edit those settings directly by right-clicking the appropriate cell.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate pages for configuring IPv4 and IPv6 access rules were unified. However, for the earlier ASA versions, a separate page for IPv6 access rules is still provided. The following descriptions apply to apply to all versions of the page, except where noted.

Navigation Path

To access the Advanced dialog box:

- In the Add and Edit Access Rule Dialog Boxes, page 16-14, click the **Advanced** button.

To access one of the Edit options dialog boxes:

- Right-click the Options or Expiration Date cell in an access rule (on the Access Rules Page, page 16-10) and choose the related Edit command. To change the rule direction, right-click the Dir. cell and choose the opposite direction (in or out).

  If you select multiple rows, your changes replace those options for all selected rules.

Related Topics

- Configuring Access Rules, page 16-7
- Editing Rules, page 12-10
- Understanding Access Rules, page 16-1
- Chapter 16, “Managing Firewall Access Rules”
- Configuring Time Range Objects, page 6-71

Field Reference

Table 16-3 Advanced Dialog Box

Element	Description
Enable Logging (PIX, ASA, FWSM)	Whether to generate syslog messages for the rule entries (also known as access-control entries, or ACEs), for PIX, ASA, and FWSM devices. When selected, these additional options are enabled: • Default Logging—Use the default logging behavior. If a packet is denied, message 106023 is generated. If a packet is permitted, no syslog message is generated. The default logging interval is 300 seconds. • Per ACE Logging—Configure logging specific to this entry. Choose the logging Level you want to use to log events for the ACE, and provide a logging Interval, which can range from 1 to 600 seconds. Syslog message 106100 is generated for the ACE. Available logging levels: – Emergency—(0) System is unstable – Alert—(1) Immediate action is needed – Critical—(2) Critical conditions – Error—(3) Error conditions – Warning—(4) Warning conditions – Notification—(5) Normal but significant condition – Informational—(6) Informational messages only – Debugging—(7) Debugging messages **Note** You can change the firewall and IOS logging options for an existing rule in the table on the Access Rules Page, page 16-10 by right-clicking the Options cell and choosing Edit Options.
Enable Logging (IOS) Log Input (IPv4 only; neither option presented on the IPv6 Access Control page)	Whether to generate an informational logging message about the packet that matches the entry; the message will be sent to the console for IOS devices. Select Log Input to include the input interface and source MAC address or virtual circuit in the logging output.
Traffic Direction	For interface-specific access rules, the direction of the traffic to which this rule applies: • In—Packets entering an interface. • Out—Packets exiting an interface. **Note** You can change the direction for an existing rule in the table on the Access Rules Page, page 16-10 by right-clicking the Dir. cell and choosing the opposite direction. Global rules are always applied in the In direction, so you cannot change this setting when configuring a global rule.

Table 16-3 Advanced Dialog Box (continued)

Element	Description
Time Range	The name of a time range policy object that defines the times when this rule applies. The time is based on the system clock of the device. The feature works best if you use NTP to configure the system clock. Enter the name or Select the object. If the object that you want is not listed, click the Create button to create it. **Note** Time range is not supported on FWSM 2.x or PIX 6.3 devices.
Options (IOS) (IPv4 only; not presented on the IPv6 Access Control page)	Additional options for IOS devices: • none—Do not apply. • Fragment—Allow fragmentation, which provides additional management of packet fragmentation and improves compatibility with NFS. By default, a maximum of 24 fragments is accepted to reconstruct a full IP packet. However, based on your network security policy, you might want to consider configuring the device to prevent fragmented packets from traversing the firewall. • Established—Allow outbound TCP connections return access through the device. This option works with two connections: an original connection outbound from a network protected by the device, and a return connection inbound between the same two devices on an external host.
Rule Expiration	Lets you configure an expiration date for the rule. Click the calendar icon to select a date. For more information, see Configuring Expiration Dates for Access Rules, page 16-22. If you configure an expiration date, you can also configure the number of days before the rule expires to send out a notification of the pending expiration, and e-mail addresses to which to send the notifications. These fields are initially filled with the information configured on the Rule Expiration administrative settings page (select **Tools > Security Manager Administration > Rule Expiration**). You can change these options for an existing rule in the table on the Access Rules Page, page 16-10 by right-clicking the Expiration Date cell and choosing Edit Rule Expiration. **Note** Expired rules are not automatically deleted. You must delete them yourself and redeploy the configuration to the device.

Hit Count Selection Summary Dialog Box

Use the Hit Count Selection Summary dialog box to select the rules for which you want to refresh hit count information. Your options are limited by the rules you selected before clicking the Refresh Hit Count button. When you click OK in this dialog box, updated hit count information is obtained from the device, which can take some time, so you are given the option to abort the operation.

Note The Hit Count of a duplicated ACE, either within the same rule or different rules, is always set to 0.

Tip You can view detailed hit count information for a rule by right-clicking the Hit Count cell for that rule on the Access Rules Page, page 16-10. Detailed hit count information is displayed in the Hit Count Details window, as described in Viewing Hit Count Details, page 16-36.

Navigation Path

(Device view only) On the Access Rules Page, page 16-10, select one access rule in the table for which you want detailed hit count information, then right-click the Hit Count column and choose Show Hit Count Details.

Related Topics

- Viewing Hit Count Details, page 16-36
- Understanding Access Rules, page 16-1

Field Reference

Table 16-4 Hit Count Selection Summary Dialog Box

Element	Description
Policy Selected	Identifies the selected policy. If you do not select a policy, this is Local, which means the rules defined specifically for the device. The policy might also be a scope within a shared or inherited policy. The indication in this field does not actually limit the scope of your hit count report.
Rules Selected	The rules for which you want to obtain hit count details; choose: • Select the **rules** option to obtain information for only those rules you selected. You can select the rows related to the name of a scope, a section name, multiple individual rules, or create a filter and select all filtered rules. This is the default if any row is selected when you initiate the hit count report. • Select **All Rules** to get hit counts for all inherited, shared, and local rules. The option is not restricted to the scope indicated in the Policy Selected field. This is the only available option if you do not select any rules before initiating the hit count report.

Table 16-4 Hit Count Selection Summary Dialog Box (continued)

Element	Description
Fetch Data From	Select one of the following options and then click **Refresh Hit Count**: • Device—Security Manager fetches the Hit Count information from the device and displays the same on the Access Rules policy page. Beginning from version 4.9, Security Manager stores the Hit Count information in its database for ASA and ASASM devices. • History—Security Manager fetches the latest Hit Count information for the particular ACEs from its database (the Hit Count history) and displays the same on the Access Rules policy page. **NOTE:** If there is an open activity the Hit Count data will not be persisted in the Security Manager database. This feature is supported in IPv4 Access Rules starting from ASASM/ASA version 8.3 and Unified Access Rules starting from ASASM/ASA version 9.0. If the Fetch Data From Device is based on Rules Selected option, the Hit Count persistence support will not be enabled. The support of Hit Count persistence will be enabled only if you choose the All Rules option within Fetch Data from Device. While fetching data From history, if the value of Hit Count is zero, Security Manager checks whether the rule was hit earlier based on the Hit Count History of the rule and displays its corresponding values. If Security Manager is not able to find the values of previous Hit from History, it displays the value of Hit Count as zero.

Configuring Expiration Dates for Access Rules

A frequent use of access rules is to provide temporary access to a network. For example, you might configure an access rule to allow a partner access for the duration of a specific project. Ideally, you want to remove the access rule at the completion of the project. However, as access rule lists grow, it is hard to manage them and to remember which rules were meant to be temporary.

To help you deal with this problem, you can configure expiration dates for access rules. By configuring an expiration date, you can project when you believe the rule will no longer be needed.

Expiration dates are not hard and fast dates; Security Manager does not delete rules that reach their expiration date. Instead, when an expiration date is reached, Security Manager displays “Expired” in bold letters in the Expiration Date column for the rule. You can filter the access rules page based on the expiration date field, for example, filtering for “expiration date has passed” to see all expired rules.

If the rule is no longer needed, you can delete it (right-click and select **Delete Row**), or disable it (right-click and select **Disable**), and then redeploy the configuration to the device. You might want to initially disable the rule, which leaves the rule in the table (overlain with hash marks), in case you discover the rule really was still needed after all, saving you the time of recreating the rule. You then just need to re-enable the rule (right-click and select **Enable**) and redeploy the configuration.

When you configure an expiration date, you can also configure notification settings, specifying an e-mail address that should be notified when an expiration date is approaching. You can specify how many days before the expiration date to send the notification e-mail message to allow you time to evaluate the rule.

The notification settings are initially filled with the values configured in the administration settings (select **Tools > Security Manager Administration > Rule Expiration**); you can enter different settings for a given rule.

To configure rule expiration:

- When creating a new rule, or editing an entire rule, click the Advanced button in the Add and Edit Access Rule Dialog Boxes, page 16-14 to get to the rule expiration settings.
- For existing rules, you can add or edit expiration settings without editing the entire rule. Right-click the Expiration Date cell for the rule and select **Edit Rule Expiration**. You can select multiple rows to configure the same rule expiration settings. For more information, see Advanced and Edit Options Dialog Boxes, page 16-17.

Related Topics

- Rule Expiration Page, page 11-69
- Configuring Access Rules, page 16-7

Configuring Settings for Access Control

You can configure various settings that apply to security-device access control lists. These settings work in conjunction with your access rules policy. The main setting of interest is that you can configure your own ACL names for each interface/traffic direction combination, or for the global ACL on ASA 8.3+ devices. For PIX, ASA, and FWSM devices, you can also control the maximum number of concurrent flows and the related syslog interval.

You can also configure an interface to allow per-user downloadable ACLs for PIX, ASA, and FWSM devices. This allows you to configure user-based ACLs in your AAA server to override the ACLs defined on a device.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate pages for configuring IPv4 and IPv6 access control were unified. However, for the earlier ASA versions, a separate page for IPv6 settings is still provided. The following descriptions apply to apply to all versions of the page, except where noted.

Related Topics

- Configuring Access Rules, page 16-7

Step 1 Do one of the following to open the Access Control Settings Page, page 16-24:

- (Device view) Select **Firewall > Settings > Access Control** (or **Firewall > Settings > IPv6 Access Control**) from the Policies selector.
- (Policy view) Select **Firewall > Settings > Access Control** (or **Firewall > Settings > IPv6 Access Control**) from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Configure the global settings in the top part of the page. For PIX, ASA, and FWSM devices, you can define the maximum number of concurrent deny flows and the related syslog interval. For ASA 8.3+ devices, you can enable object group search to optimize ACL performance when converting from Checkpoint, but this setting is not recommended unless you have a memory-constrained device. (Not available on the IPv6 Access Control page.)

For specific information about these settings, and the platforms that support ACL compilation, see Access Control Settings Page, page 16-24.

Step 3 For each interface on which you want to configure an ACL name, or enable per-user ACLs, add the interface to the interfaces table by clicking the **Add Row** button beneath the table and filling in the Firewall ACL Setting Dialog Box, page 16-26. Keep the following in mind:

- If you configure an ACL name, the name is applied to the specific interface and direction. Security Manager creates system-generated names for any interface/direction combinations that you do not specifically name.
- You can also specify the name of the global ACL for ASA 8.3+ devices.

You can edit existing entries in the list by selecting them and clicking **Edit Row**, or delete them by clicking **Delete Row**.

Access Control Settings Page

Use the Access Control Settings page to configure settings to use in conjunction with your access rules policy. You can control some performance and logging features, and configure ACL names for individual interfaces.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate policies and objects for configuring IPv4 and IPv6 access control were "unified," meaning one set of rules in which you can use either IPv4 or IPv6 addresses, or a mixture of both. However, for the earlier ASA versions, a separate page for IPv6 settings is still provided. (See Policy Object Changes in Security Manager 4.4, page 1-10 for additional information.) The following descriptions apply to apply to all versions of the page, except where noted.

Thus, many of these settings apply only to specific device types or software versions. If you configure an option and apply the policy to unsupported device types, the option is ignored for those unsupported devices.

Navigation Path

To open the Access Control Page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Access Control** (or **Firewall > Settings > IPv6 Access Control**) from the Policies selector.
- (Policy view) Select **Firewall > Settings > Access Control** (or **Firewall > Settings > IPv6 Access Control**) from the Policy Type selector. Create a new policy or select an existing policy.
- (Map view) Right-click a device and select **Edit Firewall Settings > Access Control** (or **Edit Firewall Settings > IPv6 Access Control**).

Related Topics

- Configuring Settings for Access Control, page 16-23
- Understanding Access Rules, page 16-1
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Understanding Access Rules, page 16-1

- Understanding Interface Role Objects, page 6-73

Field Reference

Table 16-5 Access Control Settings Page

Element	Description
Maximum number of concurrent flows (PIX, ASA, FWSM) (not presented on the IPv6 Access Control page)	The maximum number of concurrent deny flows that the device is allowed to create. Syslog message 106101 is generated when the device reaches the number. The range you should use depends on the amount of flash memory available in the device: • More than 64 MB—Values are 1-4096. The default is 4096. • More than 16 MB—Values are 1-1024. The default is 1024. • Less than or equal to 16 MB—Values are 1-256. The default is 256.
Syslog interval (PIX, ASA, FWSM) (not presented on the IPv6 Access Control page)	The interval of time for generating syslog message 106101, which alerts you that the security appliance has reached a deny flow maximum. When the deny flow maximum is reached, another 106101 message is generated if the specified number of seconds has passed since the last 106101 message. Values are 1-3600 milliseconds. The default is 300.
Enable Access List Compilation (Global) (IPv4 only; also not presented on the IPv6 Access Control page)	Whether to compile access lists, which speeds up the processing of large rules tables. Compilation optimizes your policy rules and performance for all ACLs, but is supported on a limited number of older platforms: • Routers (global configuration only): 7120, 7140, 7200, 7304, and 7500. • PIX 6.3 firewalls, in global mode or per interface. An ACL is compiled only if the number of access list elements is greater than or equal to 19. The maximum recommended number of entries is 16,000. To compile access lists, the device must have a minimum of 2.1 MB of memory. Access list compilation is also known as Turbo ACL.
Enable Object Group Search (ASA 8.3+) (not presented on the IPv6 Access Control page)	Whether to enable object group search on ASA 8.3+ devices, which optimizes ACL performance without expanding object groups. Object group search is mainly for use when migrating from Checkpoint to ASA, which can result in a large increase in the number of access rules, when you have a memory-constrained device (that is, you find during operations that memory runs low). If you enable object group search, you cannot use the Hit Count tool to analyze your rules. In most cases, you should not enable this feature. Instead, use the rule combination tool to simplify your access rules, and consider using global rules for rules you want to enforce on all interfaces.

Table 16-5 Access Control Settings Page (continued)

Element	Description
Access Control settings table	The table lists the interfaces for which you want to configure special processing. The interface name can be a specific interface or an interface role (which can apply settings to more than one interface at a time), or Global for global ACL settings on ASA 8.3+ devices. The main use of this table is to configure names for ACLs if you do not want Security Manager to configure system-generated names. The name applies to the ACL generated for an interface in a specific direction. You can also configure interface-level settings for per user downloadable ACLs, object group search, and ACL compilation. • To add an Access Control interface setting, click the Add Row button and fill in the Firewall ACL Setting Dialog Box, page 16-26. • To edit an Access Control interface setting, select it and click the Edit Row button. • To delete an Access Control interface setting, select it and click the Delete Row button.

Firewall ACL Setting Dialog Box

Use the Firewall ACL Setting dialog box to configure settings for specific interfaces, interface roles, or global rules for use with security-device access rules policies.

Navigation Path

Go to the Access Control Settings Page, page 16-24 and click the **Add Row** button below the interface table, or select a row in the table and click the **Edit Row** button.

Related Topics

- Configuring Settings for Access Control, page 16-23
- Understanding Access Rules, page 16-1
- Understanding Global Access Rules, page 16-3
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 16-6 Firewall ACL Setting Dialog Box

Element	Description
Interface Global (ASA 8.3+)	Specify whether you are configuring settings for specific interfaces (or interface roles), or for global rules on ASA 8.3+ devices. If you select **Interface**, specify the name of the interface or interface role for which you are configuring settings. Enter the name or click **Select** to select it from a list or to create a new object. If you select Global, your only option is to specify the name of the global ACL.
Traffic Direction	The direction of the traffic through the interface, In or Out. The settings you configure apply only to this direction, if direction matters. For global ACLs on ASA 8.3+ devices, the direction is always in.
User Defined ACL Name (checkbox not presented on the IPv6 Access Control page) ACL Name	Whether you want to supply the name for the ACL. If you select this option, enter the name you want to use, which is applied to the ACL generated for the interface and direction combination. The name must be unique on the device. If you are configuring the name for the global ACL on ASA 8.3+ devices, the option is automatically selected; simply enter the desired name. **Note** Make sure that the firewall rules ACL name is unique and is not the same name as the ACL object defined in the Policy Object Manager. For more information see Creating Access Control List Objects, page 6-53. If you do not configure a name, Security Manager generates a name for you.
Enable Per User Downloadable ACLs (PIX, ASA, FWSM) (not presented on the IPv6 Access Control page)	Whether to enable the download of per-user ACLs to override the ACLs on the interface. User ACLs are configured in a AAA server; they are not configured in Security Manager. If there are no per-user ACLs, the access rules configured for the interface are applied to the traffic. The option is configured on the device for the specified interface only when the Traffic Direction is in.
Enable Object Group Search (PIX 6.x) (not presented on the IPv6 Access Control page)	Whether to enable object group search on a PIX 6.x interface, which reduces the memory requirement on the device to hold large ACLs. However, object group search impacts performance by making ACL processing slower for each packet. Object group search is recommended when you have very large object groups. **Tip** If you are trying to configure object group search on ASA 8.3+ devices, the setting is on the Access Control Settings Page, page 16-24.

Table 16-6 Firewall ACL Setting Dialog Box (continued)

Element	Description
Enable Access List Compilation (PIX 6.x) (not presented on the IPv6 Access Control page)	Whether to compile access lists on this interface for PIX 6.x devices. This setting overrides the equivalent global setting that you configure on the Access Control Settings page. ACL compilation speeds the processing of large rules tables and optimizes your policy rules and performance for the interface. An ACL is compiled only if the number of access list elements is greater than or equal to 19. The maximum recommended number of entries is 16,000. To compile access lists, the device must have a minimum of 2.1 MB of memory.

Using Automatic Conflict Detection

Security Manager provides an Automatic Conflict Detection feature for access rules. You can use Automatic Conflict Detection to evaluate the logic of your access rules. When enabled, Automatic Conflict Detection identifies rules that overlap or conflict with other rules in the access rule policy. Use this information to identify rules that need to be deleted, moved, or edited.

This section contains the following topics:

- Understanding Automatic Conflict Detection, page 16-28
- Understanding the Automatic Conflict Detection User Interface, page 16-30
- Resolving Conflicts, page 16-34

Understanding Automatic Conflict Detection

Security Manager provides an automatic conflict detection feature to help identify unnecessary redundant or duplicate rules. Certain conflicting rules might have no effect on a device after they are deployed; however, they create unnecessary clusters in the rules table. By detecting these rules, you can clean up the rule set to develop an easier to use and more efficient access rules policy.

Other conflicting rules, can create unwanted results to your network. By detecting these conflicting rules, you can identify rules that need to be deleted, moved, or edited to implement your security needs as intended.

Note The conflict detection feature will report on the first conflict between two rules. If there are additional rules in the table that also conflict with a rule, they will not be reported until the first conflict is resolved.

Conflicts detected by Security Manager are categorized in the following way:

- Redundant Object—An element in a field of a rule is a subset of one or more elements in the same field of the rule. In the following example, the source cell has two network objects: *net-group2* and *net-group1*. Since *net-group2* is a sub-set of *net-group1*, it is a redundant object and can safely be removed:

```
object-group network net-group1
network-object 10.2.0.0 255.255.0.0

object-group network net-group2
```

```
network-object 10.2.1.1 255.255.255.255
```

- Redundant Rule—Two rules apply the same action to the same type of traffic, and removing the base rule would not change the ultimate result. For example, if a rule permitting FTP traffic for a particular network were followed by a rule allowing IP traffic for that same network, and there were no rules in between denying access, then the first rule is redundant and can be deleted.

  The following is a simple example of redundant rules:

  ```
  access-list acl permit ip 2.1.1.1 255.255.255.255 any
  access-list acl permit ip 2.1.1.0 255.255.255.0 any
  ```

- Partially Redundant Rule—A portion of a compound rule is redundant to a rule or a portion of a compound rule that follows it.
- Shadowed Rule—This is the reverse of a redundant rule. In this case, one rule will match the same traffic as another rule such that the second rule will never be applied to any traffic because it comes later in the access list. If the action for both rules are the same, you can delete the shadowed rule. If the two rules specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy. For example, the base rule might deny IP traffic, and the shadowed rule might permit FTP traffic, for a given source or destination.

  The following is a simple example of shadowed rules:

  ```
  access-list acl permit ip 1.0.0.0 255.0.0.0 any
  access-list acl permit ip 1.1.0.0 255.255.0.0 any
  ```

Note Duplicate rules are reported as shadowed rules by the automatic conflict detection feature.

- Partially Shadowed Rule—A portion of a compound rule is shadowed by a rule before it. If the action for both rules are the same, you can delete the portion of the rule that is shadowed. If the two rules specify different actions for traffic, you might need to move the shadowed rule or edit one of the two rules to implement your desired policy.

Scope of Automatic Conflict Detection

When detecting conflicts, Security Manager evaluates the following pieces of information in your access rules:

- source
- destination
- services
- users
- interfaces

Note Conflict detection is only available for access rules in the Access Rules policy for a device or shared policy. Conflict detection is not available for access rules that are part of other policies, such as AAA or inspection rules.

Note If a rule contains an FQDN network/host object, the FQDN object is ignored, but the rule is otherwise included in the analysis.

Note Disabled rules are not evaluated during conflict detection.

Note Conflict detection does not consider time ranges when evaluating access rules. Make sure that such rules truly conflict before removing any rules flagged during conflict detection.

Related Topics

- Understanding the Automatic Conflict Detection User Interface, page 16-30
- Resolving Conflicts, page 16-34
- Understanding Access Rules, page 16-1
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7

Understanding the Automatic Conflict Detection User Interface

The Automatic Conflict Detection feature is tightly coupled with the access rules table to make identifying conflicts and then resolving those conflicts faster and easier. When conflict detection is enabled, additional user interface elements are available for navigating between the conflicts and for resolving those conflicts.

Note For information on the standard elements of the Access Rules page, see Access Rules Page, page 16-10.

Figure 16-2 Automatic Conflict Detection

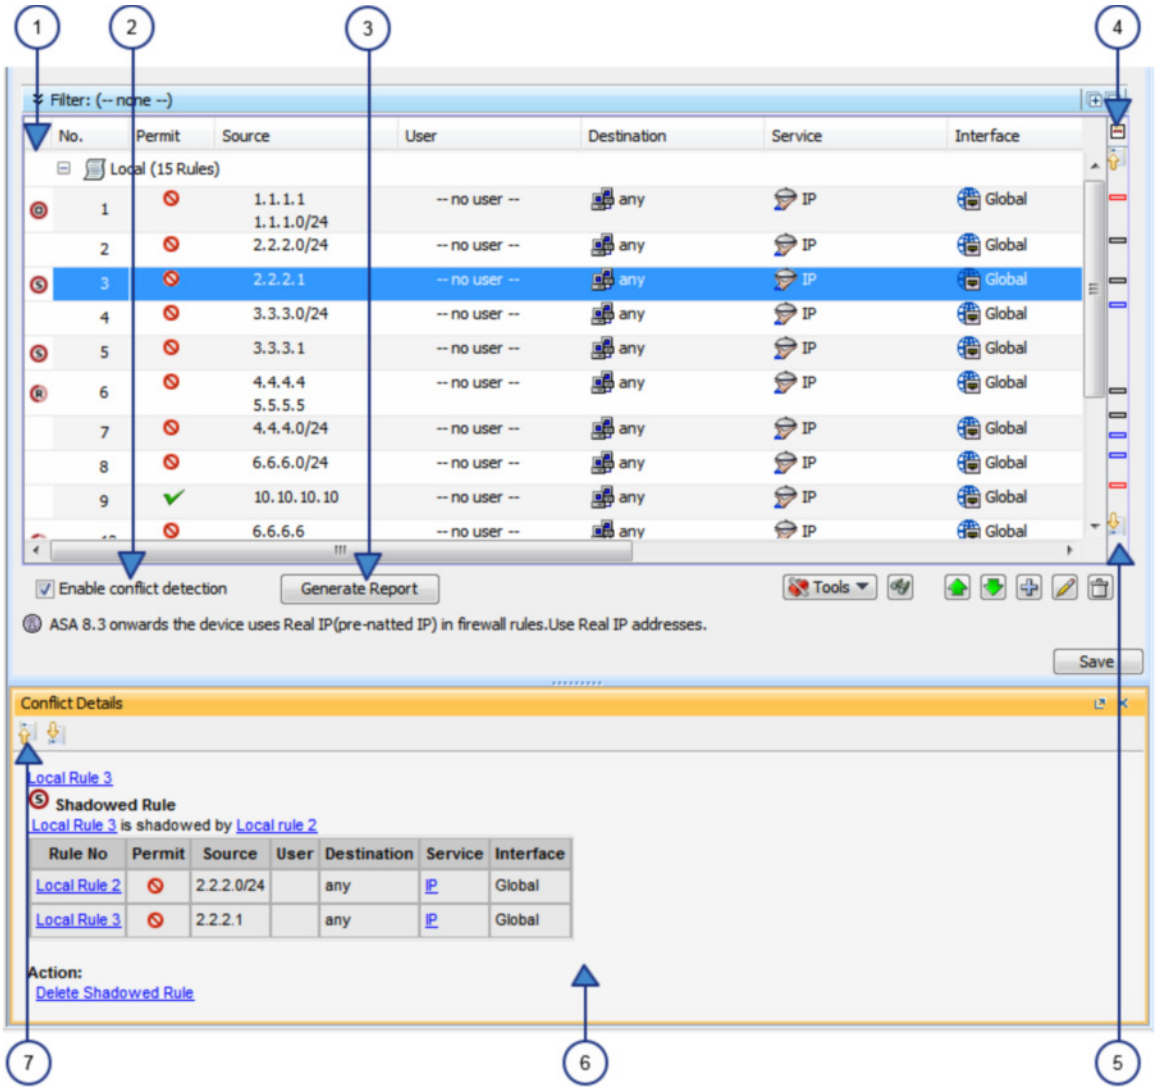

1	Conflict Indicator icons	**2**	Enable Conflict Detection
3	Generate Report button	**4**	Annotation Display Options
5	Conflict navigation bar	**6**	Conflict Details area
7	Conflict Navigation buttons		

Conflict Indicator Icons

The Conflict Indicator icons are used to identify conflicts and to provide a quick visual representation of the type of conflict. The following table details the available icons:

Note For an explanation of the types of conflicts, see Understanding Automatic Conflict Detection, page 16-28.

	Redundant Object
	Redundant Rule
	Partially Redundant Rule
	Shadowed Rule
	Partially Shadowed Rule

Note If an access rule has multiple conflicts or if it has a user note attached to it, the conflict indicator icon for that rule will have a small plus sign (+) on it.

You can perform the following actions using the Conflict Indicator icon:

- Hover the mouse pointer over the **Conflict Indicator** icon to view a description of the conflict including any user notes attached to the conflict.
- Click the **Conflict Indicator** icon, or right-click the icon and select **Show Conflict Detail**, to open the Conflict Details pane for the selected conflict.
- Right-click the **Conflict Indicator** icon for a redundant object and select **Remove Redundant Object** to remove the redundant object from a rule.
- Right-click the **Conflict Indicator** icon and select **Add User Note** to open the Add User Note dialog box for the selected conflict. You can use the Add User Note dialog box to enter a note about the conflict that can later be included in the Rule Analysis Detail Report.

  Note User notes are not saved when leaving the access rules page or after editing a rule that has a user note.

Enable Conflict Detection

The Enable Conflict Detection option controls whether automatic conflict detection is enabled. Conflict detection is enabled by default but can be disabled by deselecting this option. The setting is managed per user and enabling or disabling conflict detection for one access rules table, will also enable or disable the feature for other access rules tables.

Generate Report

If conflict detection is enabled, you can click the **Generate Report** button to create an HTML report of the conflicts that can be printed or exported to another tool. The Rule Analysis Detail Report shows details of all the conflicts in your rules table and includes any user notes that were entered for the conflicts. It does not use the settings you selected in the Annotation Display Options dialog box and does not consider the filter settings defined for the table.

Note User notes are not saved when leaving the access rules page or after editing a rule that has a user note.

When you first open the Access Rules page, the Generate Report button is replaced with a progress bar. After conflict analysis has completed, the Generate Report button becomes available along with the other conflict detection features.

Annotation Display Options button

Click the **Annotation Display Options** button to open the Annotation Display Options dialog box, which is used for selecting the types of conflicts that should be reported. For an explanation of the types of conflicts, see Understanding Automatic Conflict Detection, page 16-28.

Disabling a certain type of conflict does not remove those rules from the access rules table; it only turns off the rule conflict notification for those types of conflicts. To hide or show only conflicting rules of a certain type, you can use the table filter feature. For example, if you only wanted to see redundant and partially redundant rule conflicts, you could set up the following advanced filter:

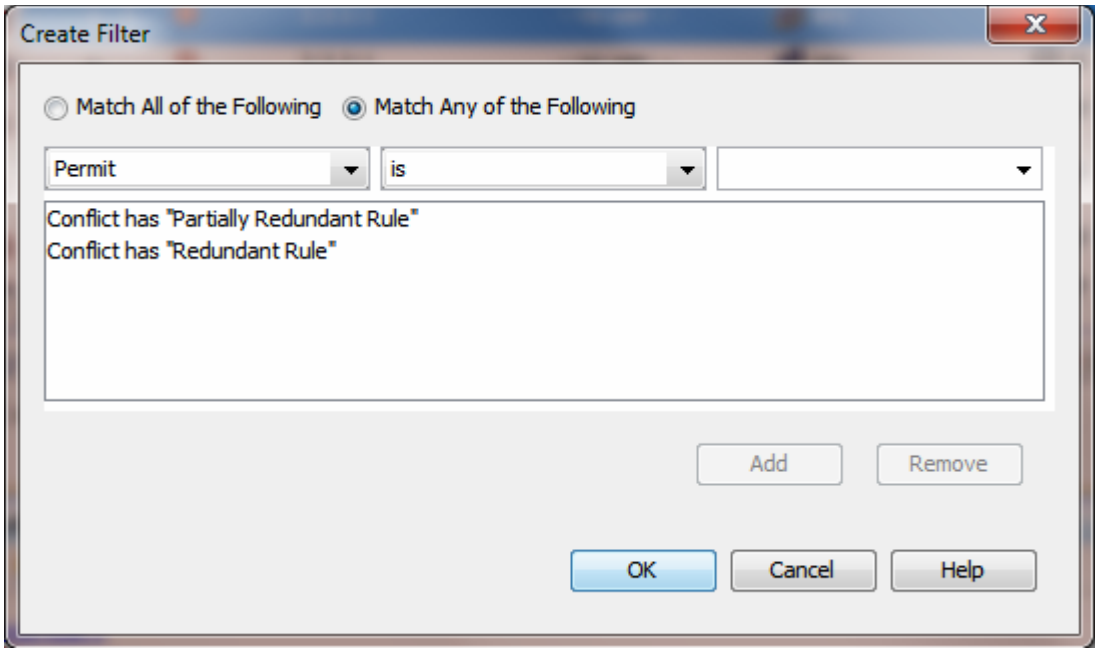

You can hover the mouse pointer over the Annotation Display Options button to view a summary of the conflicts for each type and also to see which conflict types are disabled.

Note The Annotation Display Options that you select remain in effect until those options are changed. Be sure to verify these settings whenever you are working on resolving conflicts.

Conflict Navigation Bar

Use the Conflict navigation bar to navigate to a conflict. You can use the Previous Conflict and Next Conflict buttons on the Conflict navigation bar to step through the conflicts. You can also click on one of the conflict locators in the Conflict navigation bar to move directly to a specific conflict. This is particularly helpful when working with large rules tables.

Tip Hovering over a conflict locator provides a quick summary of the conflict.

The conflict locators are color-coded as follows:

- **Red locators**—Redundant objects
- **Blue locators**—Redundant and partially redundant rules
- **Black locators**—Shadowed and partially shadowed rules

Conflict Details Area

The Conflict Details pane shows details for the selected conflict. The pane can be docked and undocked as needed. If the Conflict Details pane is docked while the Policy Object Manager pane is also docked, you can navigate between the two features using the tabs at the bottom of the window.

The conflicting rules are shown together in a table for easier direct comparison. The type of conflict is shown above the table. A suggested action is shown below the table for all conflicts except partially redundant rules and partially shadowed rules, which must be resolved manually. Links are provided for direct navigation to the rules involved. Policy objects that are part of the conflicting rules can be expanded by clicking on them to see the object contents. Click again to collapse the policy object.

You can use the links provided to navigate to the conflicting rules. You can also click the link under Action to have Security Manager perform the suggested action automatically.

Conflict Navigation Buttons

The Previous Conflict and Next Conflict buttons at the top of the Conflict Details pane allow you to step through the conflicts that need to be resolved without leaving the Conflict Details pane.

Related Topics

- Understanding Automatic Conflict Detection, page 16-28
- Resolving Conflicts, page 16-34
- Understanding Access Rules, page 16-1
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7

Resolving Conflicts

The following procedure explains how to use the Automatic Conflict Detection feature to resolve conflicts in your access rules.

Tip You can use the Combine Rules tool to have Security Manager evaluate your rules and find ways to combine them into more efficient rules. For more information, see Combining Rules, page 12-22.

Related Topics

- Understanding Automatic Conflict Detection, page 16-28
- Understanding the Automatic Conflict Detection User Interface, page 16-30
- Understanding Access Rules, page 16-1
- Understanding Device Specific Access Rule Behavior, page 16-4
- Understanding Access Rule Address Requirements and How Rules Are Deployed, page 16-5
- Configuring Access Rules, page 16-7

Step 1 Do one of the following:

- (Device view) Select **Firewall > Access Rules** from the Policy selector.

- (Policy view) Select **Firewall > Access Rules** from the Policy Type selector and select an existing policy.

This opens the Access Rules Page, page 16-10. If conflict detection is enabled, the access rules will be analyzed for conflicts after the table has been loaded. If conflict detection is not enabled, select **Enable conflict detection** to begin the conflict analysis.

The analysis progress is shown below the rules table. With the exception of the conflict detection features, you can perform functions on the rules table while the rules are being analyzed. After analysis has completed, the conflict detection features are enabled.

Step 2 Make sure that the rules you are interested in analyzing are being shown in the rules table. This includes expanding sections and making sure that if you are using filters that they are set correctly. Any rules that are being filtered or are in a section that is collapsed will not be included in the conflict detection analysis.

Tip You can use the Expand all rows/Collapse all rows buttons located in the upper-right corner of the Filter area above the access rules table, to quickly expand or collapse all sections in the rules table.

Step 3 Click the **Annotation Display Options** button, which is located above the Conflict navigation bar to the right of the vertical scroll bar, to open the Annotation Display Options dialog box. Verify that the types of conflicts you want detected are all enabled, and then click **OK**.

Tip You can hover the mouse pointer over the Annotation Display Options button to view a summary of the conflicts for each type and also to see which conflict types are disabled.

Note The Annotation Display Options that you select remain in effect until those options are changed. Be sure to verify these settings whenever you are working on resolving conflicts.

Step 4 If you would like to print or save a copy of the conflicts that are found in the rule table, click **Generate Report**.

The Rule Analysis Detail Report is opened in your browser. The Rule Analysis Detail Report shows details of all the conflicts in your rules table. It does not use the settings you selected in the Annotation Display Options dialog box and does not consider the filter settings defined for the table. You can save the report or print it as needed.

Step 5 Use the Conflict navigation bar to navigate to a conflict. You can use the Previous Conflict and Next Conflict buttons on the Conflict navigation bar to step through the conflicts. You can also click on one of the conflict locators in the Conflict navigation bar to move directly to a specific conflict. This is particularly helpful when working with large rules tables.

Tip Hovering over a conflict locator provides a quick summary of the conflict.

The conflict locators are color-coded as follows:

- **Red locators**—Redundant objects
- **Blue locators**—Redundant and partially redundant rules
- **Grey locators**—Shadowed and partially shadowed rules

Step 6 Click on the **Conflict Indicator** icon for the selected conflict to open the Conflict Details pane. For more information on the Conflict Indicator icons, see Understanding the Automatic Conflict Detection User Interface, page 16-30.

The Conflict Details pane shows details for the selected conflict. The conflicting rules are shown together in a table for easier direct comparison. The type of conflict is shown above the table. A suggested action is shown below the table for all conflicts except partially redundant rules and partially shadowed rules, which must be resolved manually. Links are provided for direct navigation to the rules involved. Policy objects that are part of the conflicting rules can be expanded by clicking on them to see the object contents. Click again to collapse the policy object.

Step 7 Use the links provided to navigate to the rules and resolve the conflict as needed or click the link under Action to have Security Manager perform the suggested action automatically.

Note If you do not want to resolve the conflict at this time, you can enter a note about the conflict by right-clicking the **Conflict Indicator** icon to the left of the conflict in the access rule table and then selecting **Add User Note**. User notes are included in the Rule Analysis Detail Report, but are not saved when leaving the access rules page or after editing a rule that has a user note.

Step 8 Use the Conflict navigation bar or the **Previous Conflict** and **Next Conflict** buttons at the top of the Conflict Details pane to access additional conflicts that need to be resolved.

Step 9 If there are any remaining conflicts that you do not want to resolve at this time, you can click **Generate Report** to print or save a copy of the remaining conflicts, if desired.

Viewing Hit Count Details

Use Hit Count Details window to view information about the number of times an access rule was applied to traffic. These rules are the ones that become interface ACLs on the device. The hit count results do not show counts for any other type of ACL (for example, those used with class maps or AAA rules).

For access rules on ASA 8.3(1) devices and later, the detailed hit count report also shows the last time the access rule policy was applied to traffic. This information is helpful for determining rules that might have been superseded by other policy changes.

Use the hit count information to help you debug your access rules. The information can help you identify rules that are never hit (which might mean you do not need them, or that they are duplicates of rules higher in the ACL), and rules that are hit often (which means you might want to refine the rules).

Tip You can click the Refresh Hit Count button at the bottom of the page to update hit count information before viewing the details for a rule. See Hit Count Selection Summary Dialog Box, page 16-20 for more information.

Consider the following points when analyzing the hit count details:

- You get best results if you deploy policies to the device before viewing hit counts. If you discover a device and then generate a hit count report before deployment, the results might be incomplete or hard to interpret. For example, an access rule might not have any hit count information.

- Hit count statistics are based on ACL, not on interface. If you select **Enable ACL Sharing for Firewall Rules** on the Security Manager Administration Deployment page (see Deployment Page, page 11-13), any shared ACL provides statistics that are combined from all interfaces that share the ACL.
- If you enable network object group optimization, as described in Optimizing Network Object Groups When Deploying Firewall Rules, page 12-35, you might not get good hit count information.
- If you enable ACL optimization, as described in Optimizing Access Rules Automatically During Deployment, page 16-46, the hit count results might have problems matching ACEs from the device to access rules. Thus, when you select an access rule, you might not get any hit count results for it.
- FQDN network/host objects are ignored. You cannot obtain hit count information on these objects.
- Hit count and last hit time information is cleared when a device is restarted.
- The Hit Count of a duplicated ACE, either within the same rule or different rules, is always set to 0.

Before You Begin

Hit count reports are subject to the following limitations:

- Hit count reports are device-specific. You can generate the report for one device at a time from Device view only. Ensure that you deploy policies to the device before generating the reports.
- If you enable object group search on an ASA 8.3+ device, you cannot use the Hit Count tool. Object group search is configured on the Access Control Settings Page, page 16-24.
- Although you can select rules that include FQDN network/host objects, the objects are ignored in the hit count results.

Navigation Path

(Device view only) From the Access Rules Page, page 16-10, right-click the Hit Count cell for a rule in the table and choose **Show Hit Count Details**.

The Hit Count Details window opens as a pane at the bottom of the access rules table. Click the expand button on the right side of its title bar to view the hit count details in a separate window.

Related Topics

- Understanding Access Rules, page 16-1
- Table Columns and Column Heading Features, page 1-49
- Using Category Objects, page 6-13

Field Reference

Table 16-7 ACE Hit Count Details Window

Element	Description
Choose	You can choose how to view the hit count information: **Expanded Table** or **Raw ACE** (both are explained below).

Table 16-7 ACE Hit Count Details Window (continued)

Element	Description
Expanded Table	This view lists hit count information for the access control list entry (ACE) for the rule selected in the Access Rules table (on the Access Rules Page, page 16-10) when you opened this window. The list contains more than one ACE if the access rule generated more than one ACE when you deployed the policy to the device. Most of the columns in this table match those of the Access Rules table; many contain the specific data configured in the ACE in place of any network/host, service, or interface role objects contained in the rule, with the exception of IOS 12.4(20)T+ devices, which show data only at the object level. Also, the name of the ACL that contains the ACE is listed. The Delta column the difference in hit count for the ACE since the last refresh. The Hit Count column shows the hits for the specific ACE rather than the overall rule. See Sample Hit Count Details Window, page 16-38 for an example of this table. **Tip** You can sort on multiple columns at the same time by pressing and holding the **Ctrl** key while you click the column headings. You can sort on all columns except Interface, Direction, and ACL Name.
Raw ACE	This view shows the actual CLI for the access control entry, along with the Hit Count and Last Hit Time. Use this information if you are more comfortable evaluating device commands. See Sample Hit Count Details Window, page 16-38 for an example of this table.

Note Beginning with version 4.9, Security Manager enables to view the Hit Count history in the Expanded Table and Raw ACE options. Click the **Show History** link to view the Hit Count history in a new window. This new Hit Count History Details window displays the Hit Count and the Last Hit Time information.

Sample Hit Count Details Window

You can generate hit count reports to determine how often each rule in your access rule policy is matched to traffic. If an access rule is deployed as multiple access control entries (ACEs), for example, when you use interface roles to define rules and the roles apply to more than one interface, you can see the separate hit count information for each ACE deployed. The hit count results do not show counts for any other type of ACL (for example, those used with class maps or AAA rules).

For access rules on ASA 8.3(1) devices and later, the hit count report also shows the last time the access rule policy was applied to traffic. This information is helpful for determining rules that might have been superseded by other policy changes.

Use the hit count information to help you debug your access rules. The information can help you identify rules that are never hit (which might mean you do not need them, or that they are duplicates of rules higher in the ACL), and rules that are hit often (which means you might want to refine the rules).

The following figures show an example of a hit count report and how to use the information.

- Figure 16-3 shows the default view. The upper table lists the rules as they exist in your access rules policy, either all rules or just the ones you selected before generating the report. When you select a rule, the ACEs created on the device for that rule are listed in the expanded table in the lower half of the window. When you initially open the report, the expanded table shows the ACEs for all policies listed in the upper table.

  Hit counts in the expanded table are for each ACE, whereas the count in the rules table is the sum of the hit counts for all ACEs created by the rule. Note that the expanded table for ASA/PIX/FWSM devices, and IOS devices lower than 12.4(20)T, shows hit counts for each element within any policy objects used in the rule, whereas for IOS 12.4(20)T+ devices, the information is only provided at the object group level.
- Figure 16-4 shows the same ACEs in CLI format. These are the ACEs as they exist in the device configuration.

For more information about how to read and interpret hit count reports, see Viewing Hit Count Details, page 16-36.

Figure 16-3 Expanded Table

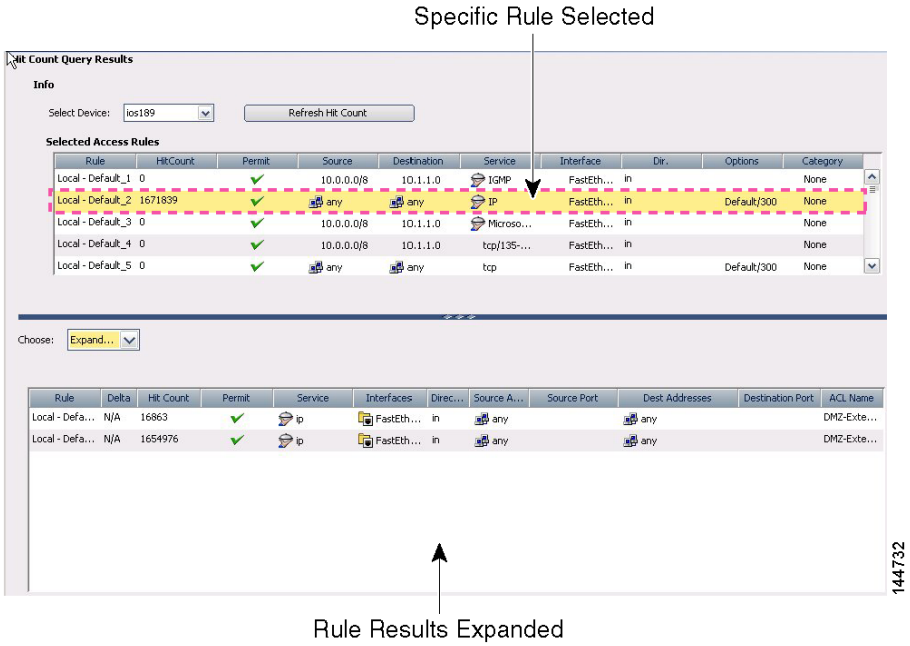

Figure 16-4 Raw ACE Table

Related Topics

- Understanding Access Rules, page 16-1
- Configuring Access Rules, page 16-7

Importing Rules

Typically, when you add a device to Security Manager, you discover policies from the device. This action populates your access rules policy with the access control entries (ACEs) from all active ACLs on the device.

If you find there are other ACLs that have ACEs you want included in your policy, you can define them directly in Security Manager.

Another alternative, however, is to import them by copying and pasting the CLI entries from a device running-configuration, or by typing in the desired commands. Using the Import Rules wizard, you can quickly create ACEs and associated policy objects from ACLs that you know already work. You might also want to use this method if you are more comfortable using CLI commands to define your rules.

The following steps describe using the Import Rules wizard to add CLI-based rules and preview the results.

Step 1 (Device view only) Select **Firewall > Access Rules** to open the Access Rules Page, page 16-10.

Step 2 Select the row after which you want to add the rules. The row must be within the Local scope. If you do not select a row, the rules are added at the end of the Local scope.

Step 3 Right-click anywhere in the rules table and choose **Import Rules** to start the wizard.

The first page—Enter Parameters—of the three-page wizard appears.

Step 4 On the Import Rules Wizard—Enter Parameters Page, page 16-41:

- Enter the desired CLI information in the *running-configuration* format appropriate for the selected device. For examples of importable CLI-based rules, see Examples of Imported Rules, page 16-44.
- Specify whether you are creating an interface-specific rule (and enter the interface or interface role to which you want the rules to apply), or for ASA 8.3+ devices, a global rule (see Understanding Global Access Rules, page 16-3).
- Specify the traffic direction with respect to the interface (the direction is always In for global rules).

Beside access control rules, you should also include the CLI information for the following items if they are referred to by the rules. If you do not include these items, the named objects must already be defined in Security Manager for the import to be successful.

- Time range objects (the **time-range** command with its subcommands), which can create time range policy objects.
- Object groups for PIX, ASA, FWSM, and IOS 12.4(20)T+ devices (the **object-group** command with its subcommands), which can create network/host policy objects.

  For ASA 8.3+ devices, you can also include the **object network** and **object service** commands. However, any object NAT configuration is not imported.

Step 5 Click **Next** to process the rules and open the Import Rules Wizard—Status Page, page 16-42.

You are notified if your CLI input contains errors when you click the Next button. For some detailed tips about what commands you can enter, see Import Rules Wizard—Enter Parameters Page, page 16-41.

The CLI is evaluated and if importable, you are told the types of objects that were created from the CLI.

Step 6 Click **Next** to view the rules and objects on the Import Rules Wizard—Preview Page, page 16-43, or click **Finish** to import the rules without previewing them.

The information on the Preview page is read-only. If the rules are acceptable, click **Finish**.

If you want to make changes, you can click the **Back** button to return to the Enter Parameters page of the wizard, or you can click Finish and edit the rules on the Access Rules page.

Import Rules Wizard—Enter Parameters Page

Use the Import Rules wizard to import a set of access control entries from an ACL in device running-configuration format to your access rules policy. The command syntax you can enter is controlled by the type of device to which you are importing rules.

Beside access control rules, you should also include the CLI for the following items if they are referred to by the rules. If you do not include these items, the named objects must already be defined in Security Manager for the import to be successful.

- Time range objects (the **time-range** command with its subcommands).
- Object groups for PIX, ASA, FWSM, and IOS 12.4(20)T devices (the **object-group** command with its subcommands).

  For ASA 8.3+ devices, you can also include **object network** and **object service** commands. However, any object NAT configuration is not imported.

Navigation Path

(Device view only) Right-click anywhere in the rules table on the Access Rules Page, page 16-10 and choose **Import Rules**.

Related Topics

- Importing Rules, page 16-40
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 16-8 Import Rules - Enter Parameters Dialog Box

Element	Description
CLI	The OS commands that define the rules and related objects that you want to import. These rules must be in running-configuration format, so they are best copied and pasted from a configuration (use Ctrl+V to paste into the field). You can also type in the commands; you will be prompted if they cannot be interpreted. You can import only one ACL at a time. To see some examples of the CLI you can import, see Examples of Imported Rules, page 16-44. **Tips** • If you refer to an object but do not include the CLI, the rule might be created but it will not use the object. • For PIX, FWSM, ASA, and IOS 12.4(20)T+, you can include object group and name commands. • If you import an ACL that is inactive, it is shown as disabled in Security Manager. If you deploy the configuration, it is removed from the device. • You can import extended ACLs for all device types, and standard ACLs for IOS devices. However, standard ACLs are converted to extended ACLs.
Interface Global (ASA 8.3+)	Select whether you are importing an interface-specific or global rule. Global rules are available only for ASA 8.3+ devices, and are handled according to special rules (for detailed information, see Understanding Global Access Rules, page 16-3). If you select **Interfaces**, enter the name of the interface or the interface role for which you are defining this rule, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. You can enter any combination of interface or interface role names, separated by commas.
Traffic Direction	The direction of the traffic with respect to the interface, in or out.
Category	The category assigned to the rules. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Import Rules Wizard—Status Page

Use the Status page of the Import Rules wizard to view information about the results of the import process.

Navigation Path

For information on starting the Import Rules wizard, see Import Rules Wizard—Enter Parameters Page, page 16-41

Related Topics

- Importing Rules, page 16-40

Field Reference

Table 16-9 Import Rules Wizard—Status Page

Element	Description
Progress bar	Shows the status of the import process.
Status	The status of the imported configuration.
Rules Imported	The number of rules that will be imported.
Policy Objects Created	The number of policy objects that will be created.
Messages	The warning, error, and informational messages, as indicated by the severity icon. Typical informational messages describe the policy objects created during the operation or the existing policy objects that were reused. When you select an item, the Description box to the right describes the message in detail. The Action box to the right provides information on how you can correct the problem.
Abort button	Click this button to stop the import operation.

Import Rules Wizard—Preview Page

Use the Preview page of the Import Rules wizard to view the rules and objects that will be imported if you click Finish.

This preview is read-only; you cannot edit the rules or objects. If the rules or objects are not exactly what you want, you can click Finish to add the rules and objects, and then edit them from the access rules page. For example, you cannot import rule expiration dates, because those dates have meaning only in Security Manager.

The tabs on this dialog box appear only if the data you are importing includes items to be displayed on the tab.

Tip If your CLI refers to an object that does not exist, such as a time range, the object is not included in the rule. You can either go back and add the CLI for the object, or you can click Finish, create the object yourself, and edit the rule.

Navigation Path

For information on starting the Import Rules wizard, see Import Rules Wizard—Enter Parameters Page, page 16-41.

Related Topics

- Importing Rules, page 16-40

- Access Rules Page, page 16-10
- Understanding Networks/Hosts Objects, page 6-80
- Understanding Interface Role Objects, page 6-73
- Understanding and Specifying Services and Service and Port List Objects, page 6-100
- Filtering Tables, page 1-48

Field Reference

Table 16-10 Import Rules Wizard—Preview Page

Element	Description
Rules tab	The rules that were created from your CLI and that will be imported to the access rules policy. All rules are converted to extended format, even if your CLI was for a standard ACL. Icons indicate the permit and deny status: • Permit—Shown as a green check mark. • Deny—Shown as a red circle with slash. You can right-click the source, destination, services, and interfaces cells and select **Show Contents** to see the detailed information in the cell. You can also right-click and select **Copy** to copy a rule to the clipboard in HTML format, which you can paste into a text editor.
Objects tab	The policy objects created from your CLI, if any. Depending on the CLI, Security Manager might create time range, network/host, service, or port list objects. Right-click an object and select **View Object** to see the object definition in read-only format.

Examples of Imported Rules

The following are some examples of CLI that you can import and the rules and policy objects that are created from them. For information on how to import rules, see Importing Rules, page 16-40.

Example 1: Restrict a network from accessing FTP servers (ASA devices)

The following access list uses object groups and restricts the 10.200.10.0/24 network from accessing some FTP servers. All other traffic is allowed.

```
object-group network ftp_servers
network-object host 172.16.56.195
network-object 192.168.1.0 255.255.255.224
access-list ACL_IN extended deny tcp 10.200.10.0 255.255.255.0 object-group ftp_servers
access-list ACL_IN extended permit ip any any
```

This example creates a network/host object named ftp_servers and two access rules.

No.	Permit	Source	Destination	Service	Interface	Dir.	Category
1		10.200.10.0/24	ftp_servers	TCP	Ethernet0	in	None
2		any	any	IP	Ethernet0	in	None

194227

Example 2: Restrict web access during working hours (ASA devices)

The following example denies HTTP requests between the hours of 8 AM and 6 PM, which are typical work hours.

```
time-range no-http
 periodic weekdays 8:00 to 18:00
access-list 101 deny tcp any any eq www time-range no-http
```

This example creates a time range object named no-http and one access rule.

No.	Permit	Source	Destination	Service	Interface	Dir.	Options	Category
1		any	any	HTTP	Ethernet0	in	no-http	None

194230

Example 3: Filtering on TCP and ICMP using port numbers (IOS devices)

In the following example, the first line of the extended access list named goodports permits any incoming TCP connections with destination ports greater than 1023. The second line permits incoming TCP connections to the Simple Mail Transfer Protocol (SMTP) port of host 172.28.1.2. The last line permits incoming ICMP messages for error feedback.

```
ip access-list extended goodports
 permit tcp any 172.28.0.0 0.0.255.255 gt 1023
 permit tcp any host 172.28.1.2 eq 25
 permit icmp any 172.28.0.0 255.255.255.255
```

This example creates three access rules. Notice that the wildcard masks used in the IOS ACL syntax are converted to regular subnet masks. Security Manager automatically converts between standard network/host subnet mask designations and the wildcard masks required in IOS ACLs. Because ASA/PIX/FWSM requires the use of subnet masks in ACL commands, this makes it possible for you to create rules that can apply to all devices; Security Manager takes care of converting your rules to the correct syntax.

No.	Per...	Source	Destination	Service	Interface	Dir.	Category
1		any	172.28.0.0/16	tcp/gt 1023	Ethernet0	in	None
2		any	172.28.1.2	SMTP	Ethernet0	in	None
3		any	any	ICMP	Ethernet0	in	None

194229

Example 4: Standard ACLs restricting hosts (IOS devices)

In the following example, the workstation belonging to Jones is allowed access to Ethernet interface 0 and the workstation belonging to Smith is not allowed access:

```
ip access-list standard workstations
 remark Permit only Jones workstation through
 permit 172.16.2.88
 remark Do not allow Smith workstation through
 deny 172.16.3.13
```

This example creates two rules, converting the standard rules to extended rules (to any destination). The remarks are saved in the description field.

For more examples of ACLs in command language format, see the following:

- IOS Devices—http://www.cisco.com/en/US/docs/ios/sec_data_plane/configuration/guide/sec_create_IP_apply.html#wp1027258.
- ASA Devices—http://www.cisco.com/en/US/docs/security/asa/asa82/configuration/guide/acl_extended.html.

Optimizing Access Rules Automatically During Deployment

You can configure the system to optimize the access control lists (ACLs) that are created from your access rules policies when they are deployed to specific devices or to all devices. This optimization affects only the deployed policies, it does not make any changes to your access rules policy.

Optimization removes redundancies and conflicts and can combine multiple entries (ACEs) into single entries. Although the order of entries might change, the semantics of your policies are preserved; the optimized ACL accepts or denies the same set of packets as did its unoptimized form. Following are the basic cases where changes are made:

- Ineffective ACE—Where one entry is a subset of, or equal to, another entry, the ineffective ACE is removed. Consider the following example:

  ```
  access-list acl_mdc_inside_access deny ip host 10.2.1.1 any
  access-list acl_mdc_inside_access deny ip 10.2.1.0 255.255.255.0 any
  ```

  The first ACE is actually a subset of the second ACE. ACL optimization will deploy only the second entry.

- Superset ACE—Where one entry is a superset of another and the order of the rules does not matter, the redundant rule is removed. Consider the following example:

  ```
  access-list acl_mdc_inside_access permit tcp any any range 110 120
  access-list acl_mdc_inside_access deny tcp any any range 115
  ```

  The second ACE will never be hit. ACL optimization will remove the second ACE and deploy only the first one.

- Adjacent ACEs—Where two entries are similar enough that a single entry can do the same job. There can be no intervening rules that change which packets will hit each rule. Consider the following example:

  ```
  access-list myacl permit ip 1.1.1.0 255.255.255.128 any
  access-list myacl permit ip 1.1.1.128 255.255.255.128 any
  ```

  The two ACEs are merged into one: access-list myacl permit ip 1.1.1.0 255.255.255.0 any.

By configuring ACL deployment optimization, you can create smaller ACLs that are more efficient, which can improve performance on devices with limited, non-expandable memory, such as the FWSM, which can be shared among multiple virtual contexts.

However, there are down sides to configuring ACL deployment optimization:

- Because optimization changes what would normally be deployed for your access rules, it is hard to correlate those rules to the actual deployed ACEs. This can make the results of the hit count tool unusable, and make it very difficult to correlate events in the Cisco Security Monitoring, Analysis and Response System application. If it is important to you that you can monitor your access rules using these tools, do not enable optimization. For more information, see Viewing Hit Count Details, page 16-36 and Viewing CS-MARS Events for an Access Rule, page 72-42.
- Optimization does not address inherent problems in your access rules policy. It is typically better to address redundancies and conflicts proactively by using the automatic conflict detection tool (see Using Automatic Conflict Detection, page 16-28). You can also use the combine rules tool to optimize your rules in the access rules policy before you deploy them (see Combining Rules, page 12-22).

If you decide to configure ACL deployment optimization, consider enabling it only for those devices that are memory constrained.

Step 1 Log into Windows on the Security Manager server.

Step 2 Use a text editor such as NotePad to open the **C:\Program Files\CSCOpx\MDC\athena\config\csm.properties** file. Locate the optimization section and read the instructions.

- To turn on full optimization for all devices, enter the following:

  OPTIMIZE.*=full

- To turn on full optimization for a specific device, replace the asterisk with the Security Manager display name for the device. For example, if the display name is west_coast.cisco.com, enter the following:

  OPTIMIZE.west_coast.cisco.com=full

- To turn on optimization but preserve the object groups used in the ACE, replace the full keyword with preserve_og. For example:

  OPTIMIZE.west_coast.cisco.com=preserve_og

- If you do not want to allow the merger of adjacent entries, enter the following:

  AclOptimization.doMerge=false

Step 3 Save the file. The settings take effect immediately and will be applied to all subsequent deployment jobs.

You can generate optimization reports for deployment jobs by selecting **Capture Discovery/Deployment Debugging Snapshots to File**, which is located in **Tools > Security Manager Administration > Debug Options**.

The deployment results will show optimization results summarized as an informational message that includes the original number of ACEs before optimization and the number of ACEs after optimization. The results are saved to a file on the server in the C:\Program Files\CSCOpx\MDC\temp folder. A job ID is used as part of the file name.

Customizing defaults in the Add Access Rule dialog

Prior to Cisco Security Manager 4.13, the **Add Access Rule** dialog was populated with default values. Starting from 4.13, the user can customize the appearance of default values by updating the **csm.properties** file.

To customize defaults in the Add Access Rule dialog, perform the following steps:

Step 1 Close and exit the Cisco Security Manager interface.

Step 2 Use a text editor such as NotePad to open the **C:\Program Files\CSCOpx\MDC\athena\config\csm.properties** file.

Step 3 Locate the CustDesk.Rule property at the bottom of the csm.properties file and set the values as true or false, based on your requirements:

- CustDesk.Rule.Add.Op.Load.Intf.Default.Values - Set this value as true to load the interface information by default in the Add Access rule dialog.
- CustDesk.Rule.Add.Op.Load.Other.Default.Values - Set this value as true to load the other values by default in the Add Access rule dialog.

Step 4 Save the file.

Note The settings do not take effect immediately. Restart the Cisco Security Manger services for the customized default values to take effect.

Step 5 Launch the Cisco Security Manager interface again.

Managing Firewall Inspection Rules

Inspection rules configure protocol inspection on a device. Inspection opens temporary holes in your access rules to allow return traffic for connections initiated within your trusted network. When traffic is inspected, the device also implements additional controls to eliminate mal-formed packets based on the inspected protocols.

Note From version 4.17, though Cisco Security Manager continues to support PIX, FWSM, and IPS features/functionality, it does not support any enhancements.

The device commands generated for inspection rules vary based on device type. For devices running ASA, PIX 7.0+, and FWSM 3.x+, access-list, policy-map, and class-map commands are used. For older FWSM and PIX 6.3 devices, fixup commands are used. For IOS devices, ip-inspect commands are used.

The following topics will help you work with inspection rules:

- Understanding Inspection Rules, page 17-2
  - Choosing the Interfaces for Inspection Rules, page 17-2
  - Selecting Which Protocols To Inspect, page 17-3
  - Understanding Access Rule Requirements for Inspection Rules, page 17-4
  - Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5
- Inspection Rules Page, page 17-8
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Settings for Inspection Rules for IOS Devices, page 17-111

The following topics can help you with general rule table usage:

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19

Understanding Inspection Rules

Inspection rules configure Context-Based Access Control (CBAC) inspection commands. CBAC inspects traffic that travels through the device to discover and manage state information for TCP and UDP sessions. The device uses this state information to create temporary openings to allow return traffic and additional data connections for permissible sessions.

CBAC creates temporary openings in access lists at firewall interfaces. These openings are created when inspected traffic exits your internal network through the firewall. The openings allow returning traffic (that would normally be blocked) and additional data channels to enter your internal network back through the firewall. The traffic is allowed back through the firewall only if it is part of the same session as the original traffic that triggered inspection when exiting through the firewall.

Inspection rules are applied after your access rules, so any traffic that you deny in the access rule is not inspected. The traffic must be allowed by the access rules at both the input and output interfaces to be inspected. Whereas access rules allow you to control connections at layer 3 (network, IP) or 4 (transport, TCP or UDP protocol), you can use inspection rules to control traffic using application-layer protocol session information.

For all protocols, when you inspect the protocol, the device provides the following functions:

- Automatically opens a return path for the traffic (reversing the source and destination addresses), so that you do not need to create an access rule to allow the return traffic. Each connection is considered a session, and the device maintains session state information and allows return traffic only for valid sessions. Protocols that use TCP contain explicit session information, whereas for UDP applications, the device models the equivalent of a session based on the source and destination addresses and the closeness in time of a sequence of UDP packets.

  These temporary access lists are created dynamically and are removed at the end of a session.
- Tracks sequence numbers in all TCP packets and drops those packets with sequence numbers that are not within expected ranges.
- Uses timeout and threshold values to manage session state information, helping to determine when to drop sessions that do not become fully established. When a session is dropped, or reset, the device informs both the source and destination of the session to reset the connection, freeing up resources and helping to mitigate potential Denial of Service (DoS) attacks.

The following topics provide more information about inspection:

- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Inspection Rules, page 17-5
- Configuring Settings for Inspection Rules for IOS Devices, page 17-111

Choosing the Interfaces for Inspection Rules

Configure inspection on devices that protect internal networks. Use it with TCP, UDP, or more specific protocols. Inspect these applications if you want the application's traffic to be permitted through the device only when the traffic session is initiated from a particular side of the device (usually from the protected internal network).

Tip For IOS devices, you need to configure inspection explicitly, and you can identify the direction of traffic to be inspected. For ASA, PIX, and FWSM devices, you cannot identify the direction, and you need to configure inspection only if you do not want the inspection defaults. In the remaining discussion, statements concerning direction apply only to IOS devices. For ASA, PIX, and FWSM, simply configure inspection on the identified interface.

In many cases, you will configure inspection in one direction only at a single interface, which causes traffic to be permitted back into the internal network only if the traffic is part of a permissible (valid, existing) session. This is a typical configuration for protecting your internal networks from traffic that originates on the Internet.

You can also configure inspection in two directions at one or more interfaces. Configure inspection in two directions when the networks on both sides of the firewall should be protected, such as with extranet or intranet configurations, and to protect against DoS attacks. For example, if the device is situated between two partner companies' networks, you might want to restrict traffic in one direction for certain applications, and restrict traffic in the opposite direction for other applications. If you are protecting a web server in the DMZ zone, you might want to configure deep inspection on HTTP traffic to identify and reset connections that have undesirable characteristics.

You might want to configure your inspection rules on the outbound interfaces of your network, those that connect to the Internet or another uncontrolled network, while allowing unfiltered connections within the trusted network. Thus, your devices use resources for inspection only on sessions that travel over unsecured and therefore potentially dangerous networks.

Related Topics

- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Inspection Rules, page 17-5

Selecting Which Protocols To Inspect

You can generically inspect TCP and UDP, which covers all applications that use these protocols. However, you can also inspect more specific protocols. In some cases, inspecting a specific protocol provides better service than generic TCP/UDP inspection. TCP and UDP inspection do not recognize application-specific commands, and therefore might not permit all return packets for an application, particularly if the return packets have a different port number than the previous exiting packet.

For example:

- Some protocols allow you to configure deep inspection. Deep inspection allows you to configure more specific rules for a traffic stream. For example, you can drop HTTP connections where the content type of the request and response do not match. For information on deep inspection and your configuration options, see Configuring Protocols and Maps for Inspection, page 17-22.
- Protocols that negotiate return channels, such as FTP, should be specifically inspected. If you use simple generic TCP inspection of FTP traffic, the negotiated channels are not opened, and the connection will fail. If you want to allow FTP, ensure that you create a specific inspection rule for it.

Multimedia protocols also negotiate return channels and should be specifically inspected. These include H.323, RTSP (Real Time Streaming Protocol), and other application-specific protocols. Some applications also use a generic TCP channel, so you might also need to configure generic TCP inspection. Any generic TCP inspection rule should appear below a more specific inspection rule in the table (that is, any rule that specifies TCP or UDP should appear at the end of the inspection rule table).

Related Topics

- Choosing the Interfaces for Inspection Rules, page 17-2
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5

Understanding Access Rule Requirements for Inspection Rules

Access rules are applied before inspection rules. Therefore, you must ensure that your access rules do not prohibit traffic that you want inspected. Use the following guidelines:

- Permit inspected traffic to leave the network through the firewall.

  All access rules that evaluate traffic leaving the protected network should permit traffic that will be inspected. For example, if Telnet will be inspected, then Telnet traffic should be permitted on all access rules that apply to traffic leaving the network.

- Deny inspected return traffic entering the network through the firewall.

  For temporary openings to be created in an access list, the access list should deny inspected return traffic because the inspection engine will open up temporary holes in the access lists for this traffic. (You want traffic to be normally blocked when it enters your network.)

- Permit or deny traffic that cannot be inspected, or that you do not want to inspect, as required by your network.

  For example, if you do not want to inspect ICMP traffic, but you want to allow some ICMP traffic, configure your access rules to allow the traffic in both directions. Consider permitting at least these ICMP message types: echo reply (for ping commands), time-exceeded (for trace route), packet-too-big (for path MTU discovery), traceroute (for trace route), and unreachable (to notify that a host cannot be found).

- Add an access rule entry denying any network traffic from a source address matching an address on the protected network.

  This is known as anti-spoofing protection because it prevents traffic from an unprotected network from assuming the identity of a device on the protected network.

- Add an entry denying broadcast messages with a source address of 255.255.255.255.

  This entry helps to prevent broadcast attacks.

Related Topics

- Understanding Access Rules, page 16-1
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Configuring Inspection Rules, page 17-5

Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX, FWSM, and IPS features/functionality, it does not support any enhancements.

Inspecting packets at the application layer, and maintaining TCP and UDP session information, provides a device with the ability to detect and prevent certain types of network attacks such as SYN-flooding. A SYN-flood attack occurs when a network attacker floods a server with a barrage of requests for connection and does not complete the connection. The resulting volume of half-open connections can overwhelm the server, causing it to deny service to valid requests. Network attacks that deny access to a network device are called denial-of-service (DoS) attacks.

Inspection helps to protect against DoS attacks in other ways. Inspection looks at packet sequence numbers in TCP connections to see if they are within expected ranges and drops any suspicious packets. You can also configure inspection to drop half-open connections, which require firewall processing and memory resources to maintain. Additionally, inspection can detect unusually high rates of new connections and issue alert messages.

For IOS devices, you can configure several inspection setting parameters to fine-tune your defenses against SYN flooding and half-open connections. Configure the **Firewall > Settings > Inspection** policy. For details about each setting, see Configuring Settings for Inspection Rules for IOS Devices, page 17-111.

Inspection can also help by protecting against certain DoS attacks involving fragmented IP packets. Even though the firewall prevents an attacker from making actual connections to a given host, the attacker can disrupt services provided by that host. This is done by sending many non-initial IP fragments or by sending complete fragmented packets through a router with an ACL that filters the first fragment of a fragmented packet. These fragments can tie up resources on the target host as it tries to reassemble the incomplete packets. To fine-tune fragment inspection, configure an inspection rule for the **fragment** protocol and configure the maximum number of fragments you want to allow and a timeout value.

Related Topics

- Understanding Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Inspection Rules, page 17-5

Configuring Inspection Rules

Inspection rules policies identify the traffic that will be inspected through an interface. Inspection tracks permitted sessions and opens temporary holes in your access rules to allow return traffic.

Inspection rules are processed after access rules, so any traffic dropped by an access rule is not inspected. You can also use deny rules to selectively exclude certain types of traffic from inspection. For example, you might create a deny inspection rule to prevent a specific class of DNS traffic from being inspected, while all other DNS traffic is inspected. The basic procedure is:

- Add a new deny rule before the default inspection rule for the specific protocol. For the Match Traffic By option, select Source and Destination Address and Port. Next, define the specific type of traffic by providing Source and Destination Network IP addresses, and selecting the desired Service type (for example, DNS-TCP). Finally, in the third screen of the inspection-rule wizard, select the appropriate protocol (for example, DNS).
- Now edit the default inspection rule (below your new deny rule in the table). Again select Source and Destination Address and Port for the Match Traffic By option. Be sure this is a Permit rule, provide an all-addresses option as the source and destination addresses, and enter IP as the Service type. In the third screen, keep the selected protocol; configure or remove the related map, as necessary.

See Inspection Rules Page, page 17-8 and Add or Edit Inspect/Application FW Rule Wizard, page 17-11 for additional information about this process.

See the following topics for more information about things you should consider when creating inspection rules:

- Understanding Inspection Rules, page 17-2
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Protocols and Maps for Inspection, page 17-22
- Understanding Map Objects, page 6-78

Before You Begin

You might have a set of inspection rules that you want to apply to all devices. To do this, you can create a shared rule and inherit its rules to each device's inspection rules policy. For more information, see Creating a New Shared Policy, page 5-54 and Inheriting or Uninheriting Rules, page 5-47.

Step 1 Do one of the following to open the Inspection Rules Page, page 17-8:

- (Device view) Select **Firewall > Inspection Rules** from the Policy selector.
- (Policy view) Select **Firewall > Inspection Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add or Edit Inspect/Application FW Rule Wizard, page 17-11.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10.

Step 3 Select whether to apply the rule to all interfaces on the device or to only the interfaces you specify.

If you elect to specify interfaces, enter the interface name or interface role, or click **Select** to select it from a list. For IOS devices, you also can select whether the rule applies in the Out direction (traffic leaving the interface). Use the In direction for all other device types.

Step 4 Select the criteria you want to use for matching traffic. This determines what gets inspected based on this rule.

- **Default Protocol Ports**—Select this option if the protocol you are inspecting uses the default ports on your network.

  If you want to constrain the inspection based on the source or destination address, also select **Limit inspection between source and destination IP addresses** (available only for ASA, PIX 7.x+, and FWSM 3.x+ devices). When you click **Next**, you are prompted for the source and destination addresses. You can specify **any** for source or destination if you are interested only in configuring the other value.

- **Custom Destination Ports**—Select this option if you want to associate additional non-default TCP or UDP ports with a given protocol, for example, treating TCP traffic on destination port 8080 as HTTP traffic. When you click **Next**, you are prompted for the port or port range.

- **Destination Address and Port (IOS devices only)**—Select this option if you want to associate additional non-default TCP or UDP ports with a given protocol only when the traffic is going to certain destinations, for example, if you want to treat TCP traffic on destination port 8080 as HTTP only when the traffic is going to server 192.168.1.10. When you click **Next**, you are prompted for the destination address and the port information.

- **Source and Destination Address and Port (PIX 7.x+, ASA, FWSM 3.x+)**—Select this option for the same reason you would select Destination Address and Port for IOS devices, although you have the additional option of identifying the source of the traffic. When you click **Next**, you are prompted for the source and destination addresses and the service port information.

Note For FWSM 2.x and PIX 6.3(x), you can select either Default Inspection Traffic or Custom Destination Ports only.

Step 5 Click **Next**. If you selected anything other than Default Protocol Ports, fill in the required addressing and port information explained above and click **Next**. See Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13.

Step 6 On the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select the protocol you want to inspect from the list. Ensure that the Device Type field indicates that inspection is supported for that protocol on the devices to which you are assigning the rule. (If you assign a rule to an unsupported device type, the rule is ignored but you will get a validation warning).

If the protocol you select allows additional configuration, the **Configure** button becomes active. Click it to view and select your options. For more information, see Configuring Protocols and Maps for Inspection, page 17-22.

For IOS devices only:

- If you selected **Custom Destination Ports** or **Destination Address and Port** as the traffic match, you can select **custom protocol** as the protocol name and click **Configure** to assign a name to the configuration.
- You can configure additional alert, audit, and timeout settings that override those set in the inspection settings policy. You can also specify whether to inspect router generated traffic for a limited number of protocols. For more information about inspection settings, see Configuring Settings for Inspection Rules for IOS Devices, page 17-111.

Step 7 Click **Finish** to save the rule.

Step 8 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.

Note From ASA 9.9.1, for cluster mode devices which are enabled with Security Gateway feature, the following list of centralized inspections are disabled:

- DCERPC
- NetBIOS
- PPTP
- RADIUS
- RSH
- SUNRPC
- TFTP
- XDMCP

During preview config, if inspection rules are configured for the unsupported devices, a validation error is displayed.

Note When there is rollback of device, the default dns policy-map configuration is automatically added to the device. Thus, after Cisco Security Manager processes the device rollback, on re-discovery of the device, the default dns-policy-map configuration is discovered in Cisco Security Manager.

Inspection Rules Page

Use the Inspection Rules page to configure inspection rules for device interfaces. Inspection examines traffic that travels through the device to discover and manage state information for TCP and UDP sessions. The device uses this state information to create temporary openings to allow return traffic and additional data connections for permissible sessions.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate policies and objects for configuring IPv4 and IPv6 inspection rules were "unified," meaning one set of inspection rules in which you can use either IPv4 or IPv6 addresses, or a mixture of both. (See Policy Object Changes in Security Manager 4.4, page 1-10 for additional information.) In Policy view, IPv4 and unified versions of the inspection policy type are provided. In addition, a utility that you can use to convert existing IPv4 policies is provided (see Converting IPv4 Rules to Unified Rules, page 12-28). The following descriptions apply to all versions of the inspection rule table, except where noted.

If you assign an IPv4 inspection-rule shared policy to a 9.0+ device, you will no longer be able to assign unified versions of those policies to that device. Likewise, if you assign a unified inspection-rule shared policy to a 9.0+ device, you will no longer be able to assign IPv4 versions of those shared policies to that device--the device will not be included in the list of available devices on the Assignments tab for the shared policy.

Inspection rules are processed after your access rules. Thus, any traffic denied by an access rule is never inspected.

Read the following topics before you configure inspection rules:

- Understanding Inspection Rules, page 17-2
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5

Tip Disabled rules are shown with hash marks covering the table row. When you deploy the configuration, disabled rules are removed from the device. For more information, see Enabling and Disabling Rules, page 12-20.

Navigation Path

To access the Inspection Rules page, do one of the following:

- (Device view) Select a device, then select **Firewall > Inspection Rules** from the Policy selector.
- (Policy view) Select **Firewall > Inspection Rules** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Policies > Inspection Rules**.

Related Topics

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19
- Using Sections to Organize Rules Tables, page 12-20
- Using Rules Tables, page 12-8
- Filtering Tables, page 1-48

Field Reference

Table 17-1 Inspection Rules Page

Element	Description
Expand all rows/Collapse all rows	Use these buttons to expand or collapse all sections in the rules table. **Note** The buttons are located in the upper-right corner of the Filter area above the inspection rules table.
Conflict Indicator icons	Identifies conflicts and provides a quick visual representation of the type of conflict. For more details, including types of conflicts and the actions you can take from this column, see Understanding the Automatic Conflict Detection User Interface, page 16-30.
No.	The ordered rule number.

Table 17-1 Inspection Rules Page (continued)

Element	Description
Permit	Whether a rule identifies traffic that should be inspected based on the conditions set: • Permit—Identifies traffic that will be inspected. Shown as a green check mark. • Deny—Exempts the traffic from inspection. Your access rules will determine if the traffic is allowed or blocked. Shown as a red circle with slash.
Sources	The sources of traffic for this rule; can be networks, security groups (ASA 9.0+ only), and users. Multiple entries are displayed on separate lines within the table cell.
Destinations	The destinations for this rule; can be networks and security groups (ASA 9.0+ only). Multiple entries are displayed on separate lines within the table cell.
Traffic Match	The type of matching used in the rule: • **default-inspection**—The rule inspects traffic based on the default port. • TCP,UDP/port number—The rule inspects traffic based on a custom port number. • Service—The rule inspects traffic based on a service specification or service object. Multiple entries are displayed as separate subfields within the table cell. See Understanding and Specifying Services and Service and Port List Objects, page 6-100.
Interface	The interfaces or interface roles to which the rule is assigned. **Global** indicates that the rule is assigned to all interfaces. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. Multiple entries are displayed as separate subfields within the table cell. See Understanding Interface Role Objects, page 6-73.
Dir.	The direction of the traffic to which this rule applies: • In—Packets entering the interface. • Out—Packets exiting the interface.
Inspected Protocol	The protocol to be inspected and possibly some configuration settings for the protocol. You can right-click this cell and choose Edit Inspected Protocol to edit this; see Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17 for more information.
Time Range	The time range policy object assigned to the rule. This object defines the time window within which inspection occurs.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of the rule, if any.

Table 17-1 Inspection Rules Page (continued)

Element	Description
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Page elements below the rules table	
Query	Click this button to run a policy query, which can help you evaluate your rules and identify ineffective rules. See Generating Policy Query Reports, page 12-28
Find and Replace button (binoculars icon)	Click this button to search for various types of items within the table and to optionally replace them. See Finding and Replacing Items in Rules Tables, page 12-16.
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope or section. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Add Row button	Click this button to add a rule to the table after the selected row using the Add or Edit Inspect/Application FW Rule Wizard, page 17-11. If you do not select a row, the rule is added at the end of the local scope. For more information about adding rules, see Adding and Removing Rules, page 12-9.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells. For more information, see Editing Rules, page 12-10.
Delete Row button	Click this button to delete the selected rule.

Add or Edit Inspect/Application FW Rule Wizard

Use the Add or Edit Inspect/Application FW Rule wizard to add and edit inspection rules. The wizard steps you through the process of configuring an inspection rule based on your selection in the **Match Traffic By** group on this page.

Read the following topics before you configure inspection rules:

- Understanding Inspection Rules, page 17-2
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5

Navigation Path

From the Inspection Rules Page, page 17-8, click the **Add Row** button or select a row and click the **Edit Row** button.

Related Topics

- Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13
- Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17
- Understanding Interface Role Objects, page 6-73
- Editing Rules, page 12-10

Field Reference

Table 17-2 Add and Edit Inspect/Application FW Rule Wizard Step 1: Traffic Match Method

Element	Description
Enable Rule	Whether to enable the rule, which means the rule becomes active when you deploy the configuration to the device. Disabled rules are shown overlain with hash marks in the rule table. For more information, see Enabling and Disabling Rules, page 12-20.
Apply the Rule to	The interface to which the rule applies: • All Interfaces—Apply the rule to all interfaces. The rule becomes a global rule on ASA, PIX, and FWSM devices. For IOS devices, the rule is configured for each interface in the In direction. • Interface (PIX 7.x+, ASA, FWSM 3.x+, IOS)—Apply the rule only to those interfaces identified in the Interfaces field. Enter the name of the interface or the interface role, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. For IOS devices only, you can select the direction of the traffic to which this rule applies, either traffic entering an interface (**In**) or exiting it (**Out**). For other devices, leave **In** as the direction.
Match Traffic By How you want to identify the traffic to inspect. If you select something other than Default Protocol Ports (by itself), you are prompted for the other port or address information when you click **Next**.	
Default Protocol Ports Limit inspection between source and destination IP addresses (PIX 7.x+, ASA, FWSM 3.x+)	Inspect traffic based on the default ports assigned to a protocol. You will select a protocol on the next page (Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17). You can also select **Limit inspection between source and destination IP addresses** to configure the inspection to occur only between a specified source and destination. Do not select this option if you want to inspect a protocol without applying any constraints to the inspected traffic. If you also select this option, the next page of the wizard is described in Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13.

Table 17-2 Add and Edit Inspect/Application FW Rule Wizard Step 1: Traffic Match Method

Element	Description
Custom Destination Ports	Inspect traffic based on specified non-default TCP or UDP destination ports. Select this option if you want to associate additional TCP or UDP traffic with a given protocol, for example, treating TCP traffic on destination port 8080 as HTTP traffic. You will specify the protocol and port(s) on the next page of the wizard; see Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13.
Destination Address and Port (IOS devices only)	Inspect traffic on IOS devices based on destination IP address and port. Select this option if you want to associate additional non-default TCP or UDP ports with a given protocol only when the traffic is going to certain destinations, for example, if you want to treat TCP traffic on destination port 8080 as HTTP only when the traffic is going to server 192.168.1.10.
Source and Destination Address and Port (PIX 7.x, ASA, FWSM 3.x)	Inspect traffic on PIX 7.x+, ASA, and FWSM 3.x+ devices based on source and destination IP addresses and services. Select this option for the same reason you would select Destination Address and Port for IOS devices, although you have the additional option of identifying the source of the traffic. You will specify the action, sources, destinations, and Services on the next page of the wizard; see Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	An optional description of the rule (up to 1024 characters).

Add or Edit Inspect/Application FW Rule Wizard, Step 2

The options presented on the second page of the Inspect/Application FW Rule Wizard depend on your **Match Traffic By** selection on the first page (see Add or Edit Inspect/Application FW Rule Wizard, page 17-11). The possibilities are as follows:

- If you select Default Protocol Ports on the first page and *do not* select Limit inspection between source and destination IP addresses, the second page consists of the options described in Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17.
- If you select Default Protocol Ports on the first page and *do* select Limit inspection between source and destination IP addresses, the second page consists of the options described in the second table in this section. (The third page will consist of the options described in Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17.)
- If you select Custom Destination Ports on the first page, the second page consists of the options described in the first table in this section. (The third page will consist of the options described in Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17.)
- If you select Source and Destination Address and Port on the first page, the second page consists of the options described in the second table in this section. (The third page will consist of the options described in Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17.)

Navigation Path

From the Add or Edit Inspect/Application FW Rule Wizard, page 17-11, select a Match Traffic By option and click Next.

Related Topics

- Understanding Inspection Rules, page 17-2
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3
- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5
- Understanding Interface Role Objects, page 6-73
- Editing Rules, page 12-10

Field Reference

The following table describes the options presented on page 2 of the Inspect/Application FW Rule Wizard after you have selected **Custom Destination Ports** on the first page of the wizard (described in Add or Edit Inspect/Application FW Rule Wizard, page 17-11).

Table 17-3 Add and Edit Inspect/Application FW Rule Wizard Step 2: Protocol and Port Page

Element	Description
Protocol	The protocol for the ports you are specifying, either TCP, UDP, or both TCP/UDP. When configuring Custom Destination Ports for an IOS device, you must select TCP/UDP.
Ports	The port(s) used by the traffic you want to inspect. Valid values range from 1 to 65535. • Single—Specify one port number only. • Range—Specify a range of ports, for example, 10000-11000. When configuring custom ports, be aware that port ranges might not be supported on all platforms or OS versions. Any conflicts are identified during policy validation, not while you are editing this rule. **Tip** If you specify a port or port range that conflicts with a pre-defined port mapping, the device does not allow the port to be remapped.

The following table describes the options presented on page 2 of the Inspect/Application FW Rule Wizard after you have selected **Default Protocol Ports** and **Limit inspection between source and destination IP addresses** on the first page of the wizard, and when you select **Source and Destination Address and Port** on the first page. The first page of the wizard is described in Add or Edit Inspect/Application FW Rule Wizard, page 17-11.

Table 17-4 Add and Edit Inspect/Application FW Rule Wizard Step 2: Action, Sources, Destinations, and Services Page

Element	Description
Action	Whether you are identifying traffic that should be inspected based on the conditions set. Typically, you will create Permit rules. • Permit—Identifies traffic that will be inspected. • Deny—Exempts the traffic from inspection. Your access rules will determine if the traffic is allowed or blocked.

Table 17-4 Add and Edit Inspect/Application FW Rule Wizard Step 2: Action, Sources, Destinations, and Services Page (continued)

Element	Description
Sources	Provide traffic sources for this rule; can be networks, security groups, and users. You can enter values or object names, or Select objects, for one or more of the following types of sources: • Network – You can specify a various network, host and interface definitions, either individually or as objects. If you Select an interface object as a source, the dialog box displays tabs to differentiate between hosts/networks and interfaces. The "All-Address" objects do not restrict the rule to specific hosts, networks, or interfaces. These addresses are IPv4 or IPv6 addresses for hosts or networks, network/host objects, interfaces, or interface roles. **Note** You can only specify a fully qualified domain name (FQDN) by providing an FQDN network/host object, or a group object that includes an FQDN object. You cannot directly type in an FQDN. See Understanding Networks/Hosts Objects, page 6-80, Specifying IP Addresses During Policy Definition, page 6-87 and Understanding Interface Role Objects, page 6-73 for additional information about these definitions. • Security Groups (ASA 9.0+) – Enter or Select the name or tag number for one or more source security groups for the rule, if any. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. • Users – Enter or Select the Active Directory (AD) user names, user groups, or identity user group objects for the rule, if any. You can enter any combination of the following: – Individual user names: NetBIOS_DOMAIN\username – User groups (note the double \): NetBIOS_DOMAIN\\user_group – Identity user group object names. For more information, see: – Selecting Identity Users in Policies, page 13-21 – Configuring Identity-Based Firewall Rules, page 13-21 – Creating Identity User Group Objects, page 13-19 **Note** Enter more than one value in any of these fields by separating the items with commas. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.

Table 17-4 Add and Edit Inspect/Application FW Rule Wizard Step 2: Action, Sources, Destinations, and Services Page (continued)

Element	Description
Destinations	Provide traffic destinations for this rule; can be networks or security groups. As with Sources, you can enter values or object names, or Select objects, for one or more destinations of Network and Security Group (ASA 9.0+) type.
Services	The services that define the type of traffic upon which to act. You can enter or Select any combination of service objects and service types (which are typically a protocol and port combination). Enter more than one value by separating the items with commas. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.
Time Range	The name of a time range policy object that defines the times when this rule applies. The time is based on the system clock of the device. The feature works best if you use NTP to configure the system clock. Enter the name or click **Select** to select the object. If the object that you want is not listed, click the **Create** button to create it.

Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page

Use the Inspect/Application FW Rule wizard's inspected protocol page to configure the protocol monitored by this inspection rule.

The options in this section are presented when you add or edit a firewall inspection rule, and when you right-click the Inspected Protocol cell of an existing rule in the table on the Inspection Rules Page, page 17-8.

Note Beginning with version 4.9, Security Manager supports SIP protocol for ASA cluster devices running the software version 9.4.0 or higher.

Navigation Path

Do one of the following:

- In the Add or Edit Inspect/Application FW Rule Wizard, page 17-11, click Next until you reach this page.
- To open the Edit Inspected Protocols dialog box, right-click the Inspected Protocol cell of an inspection rule and choose **Edit Inspected Protocol**. If you select multiple rows, your changes replace the inspected protocol defined for all selected rules.

Related Topics

- Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13
- Understanding Inspection Rules, page 17-2
- Choosing the Interfaces for Inspection Rules, page 17-2
- Selecting Which Protocols To Inspect, page 17-3

- Understanding Access Rule Requirements for Inspection Rules, page 17-4
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5
- Editing Rules, page 12-10
- Filtering Tables, page 1-48

Field Reference

Table 17-5 Inspected Protocol Options

Element	Description
Protocols table	Lists the protocols that can be inspected. You can select one protocol per rule. The list includes information on the device operating systems that allow inspection of the protocol: do not select protocols that are not supported by the device type to which you will apply the inspection rule. **Tip** For IOS devices, if you selected Custom Destination Ports or Destination Address and Port for the match type on the first page of the wizard, you can select **custom protocol** and click Configure to give your protocol a name. For other device types, select the protocol that you associate with the ports previously specified. The Options column displays configured options for the selected protocol, if any. The Group column provides additional information on the use of some of the protocols.
Selected Protocol Configure button	Displays the protocol you selected. If the protocol allows additional configuration, the Configure button becomes active; click it to see your options, and click the Help button in the dialog box that is opened for information about the options. For more information about protocols that allow configuration, see Configuring Protocols and Maps for Inspection, page 17-22.

Table 17-5 Inspected Protocol Options (continued)

Element	Description
Rule Settings (IOS)	Additional settings for the rule if it is used on devices running Cisco IOS software. If you select Use Default Inspection settings, the IOS defaults, or the settings defined in the inspection settings policy (see Configuring Settings for Inspection Rules for IOS Devices, page 17-111), are used. These are the settings you can enable or disable: • Alert—Whether to generate stateful packet inspection alert messages on the console. • Audit—Whether audit trail messages are logged to the syslog server or router. • Timeout—Whether to configure the length of time, in seconds, for which a session is managed while there is no activity. If you select Specify Timeout, enter the timeout value; the range is 5 to 43200 seconds. • Inspect Router Generated Traffic—Whether to inspect traffic that is generated by the device itself. This option is available for a limited number of the protocols.

Configure DNS Dialog Box

Use the Configure DNS dialog box to configure settings for DNS inspection on PIX 7.0+, ASA, FWSM, and IOS devices.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select DNS in the protocols table, and click **Configure**.

Field Reference

Table 17-6 Configure DNS Dialog Box

Element	Description
Maximum DNS Packet Length	The maximum DNS packet length. Values are 512 to 65535.
DNS Map	The DNS policy map object that defines traffic match conditions and actions, protocol conformance policies, and filter settings. Enter the object name, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Enable Dynamic Filter Snooping	Whether to allow the security appliance to snoop DNS packets in order to build a database of DNS lookup information. This information is used by botnet traffic filtering to match DNS names to IP addresses. If you configure a botnet traffic filtering rules policy, select this option. Otherwise, do not select the option. For more information, see Botnet Traffic Filter Rules Page, page 19-9.

Configure SMTP Dialog Box

Use the SMTP dialog box to edit settings for Simple Mail Transfer Protocol (SMTP) inspection. SMTP is used to transfer email between servers and clients on the Internet.

SMTP inspection drops any packets with illegal commands. You can configure a maximum data length for packets. Enter a length in the range 0-4294967295.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select SMTP in the protocols table, and click **Configure**.

Configure ESMTP Dialog Box

Use the Configure ESMTP dialog box to edit settings for Extended Simple Mail Transport Protocol (ESMTP) inspection. You can configure these settings based on platform:

- IOS devices—You can configure a maximum data length for packets. Enter a length in the range 0-4294967295.
- ASA/PIX 7.x+ devices—You can specify an ESMTP policy map object to define deep inspection parameters. Enter the name of the object or click **Select** to select it from a list or to create a new object.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select ESMTP in the protocols table, and click **Configure**.

Configure Fragments Dialog Box

Use the Configure Fragments dialog box to edit settings for fragment inspection on IOS devices.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select fragment in the protocols table, and click **Configure**.

Field Reference

Table 17-7 Configure Fragments Dialog Box

Element	Description
Maximum Fragments	The maximum number of unassembled packets for which state information (structures) is allocated by Cisco IOS software. Unassembled packets are packets that arrive at the router interface before the initial packet for a session. Values are 0-10000 state entries. The default is 256. **Note** Memory is allocated for the state structures, and setting this value to a larger number may cause memory resources to be exhausted.

Table 17-7 Configure Fragments Dialog Box (continued)

Element	Description
Timeout (sec)	The number of seconds that a packet state structure remains active. When the timeout value expires, the router drops the unassembled packet, freeing that structure for use by another packet. Values are 1-1000. The default timeout value is one second.

Configure IMAP or POP3 Dialog Boxes

Use the Configure IMAP or POP3 dialog boxes to edit settings for Internet Message Access Protocol (IMAP) or Post Office Protocol 3 (POP3) inspection on IOS devices.

- IMAP is a method for accessing electronic mail or bulletin board messages that are kept on a mail server that may be shared. It permits a client email program to access remote messages as though they were local.
- POP3 is used to receive email that is stored on a mail server. Unlike IMAP, POP retrieves mail only from a remote host.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select IMAP or POP3, and click **Configure**.

Field Reference

Table 17-8 Configure IMAP or POP3 Dialog Boxes

Element	Description
Reset Connection on Invalid IMAP/POP3 packet	Whether to reset, or drop, the connection between the client and server if an invalid packet is encountered. The client will have to repeat the validation process to reconnect to the server.
Enforce Secure Authentication	Whether to require that the client use a secure login to the server, that is, so that passwords are not sent in clear text.

Configure RPC Dialog Box

Use the RPC dialog box to edit settings for RPC inspection on IOS devices. RPC inspection blocks traffic for all RPC programs except for those you specify. To allow more than one RPC program, create a rule for each program number you want to allow.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select RPC in the protocols table, and click **Configure**.

Field Reference

Table 17-9 Configure RPC Dialog Box

Element	Description
Program Number	The program number to permit. Values are 1-4294967295.

Table 17-9 Configure RPC Dialog Box (continued)

Element	Description
Wait Time	The number of minutes to keep a hole in the firewall open to allow subsequent connections from the same source address to the same destination address and port. Values are 0-35791 minutes. The default is 0.

Custom Protocol Dialog Box

Use the Custom Protocol dialog box to assign a name to the protocol and port specification you made on the Add or Edit Inspect/Application FW Rule Wizard, Step 2, page 17-13 for IOS devices.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select custom protocol in the protocols table, and click **Configure**.

Configure Dialog Box

Use the Configure dialog box to select a policy map object for HTTP or IM inspection. The maps used for these types of inspection differ depending on the operating system version used on the device. Select the desired version and then click **Select** to select the desired policy map object or to create a new one.

Navigation Path

Go to the Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17, select HTTP or IM in the protocols table, and click **Configure**.

Configuring Protocols and Maps for Inspection

When you configure inspection rules for a device, you select the protocols that you want to inspect. Some of these protocols allow additional configuration for deep inspection. Deep inspection allows you to specify additional requirements that packets must meet in order to traverse the device. For example, you can drop HTTP connections where the content type of the request and response do not match. (For a full list of inspectible protocols, click **Add Row** on the Inspection Rule page and click Next to view the protocols list.)

What you can configure depends not only on the protocol but on the device's operating system and version number. Typically, your ability to fine-tune inspection is higher on an ASA device compared to an IOS device. (If you are configuring an IOS device and you want greater control over inspection, consider configuring zone-based firewall inspection; for more information, see Understanding the Zone-based Firewall Rules, page 21-3.)

Some deep inspection configuration is done directly in the inspection rule. However, for some protocols, you can configure the inspection rule to include a policy map that you create as an independent policy object. (You need to configure policy maps only if you want something other than the default inspection options.) You can configure these maps from the policy object selector dialog box while configuring the policy, or from the Policy Object Manager window (select **Manage > Policy Objects**).

For protocols that use policy maps, you can select the desired policy map, which defines the match conditions for the targeted traffic. For ASA, PIX, and FWSM devices, these policy maps might point to class maps that define the match conditions. To create these policy maps in the Policy Object Manager,

select one of the maps listed in the following table in the **Maps > Policy Maps > Inspect** folder and review the detailed usage information in the references mentioned. For information on creating class maps, which are in the **Maps > Class Maps > Inspect** folder, see the references to the match criterion dialog boxes and Configuring Class Maps for Inspection Policies, page 17-28.

Table 17-10 Configuring Protocols for Deep Inspection in Inspection Rules

Protocol	Device Types	Policy Map	Class Map (ASA, PIX, FWSM only)	Description and Match Criteria Reference
DNS	ASA, PIX, FWSM, IOS	DNS	DNS	Inspect traffic based on a wide variety of criteria using the class and policy map, which allow extensive control over DNS packets. In addition, you can configure a maximum length in the inspection rule, and enable dynamic DNS snooping for use with Botnet rules (on ASA devices). See the following topics: • Configuring DNS Maps, page 17-32 • DNS Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-36 • Configure DNS Dialog Box, page 17-19
FTP Strict	ASA, PIX, FWSM, IOS	FTP	FTP	Inspect traffic based on file name, type, server, user, or FTP command. See Configuring FTP Maps, page 17-42 and FTP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-43.
GTP	ASA, PIX, FWSM, IOS	GTP	GTP	Inspect traffic based on timeout values, message sizes, tunnel counts, and GTP versions traversing the security appliance. See Configuring GTP Maps, page 17-45 and GTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-49.

Table 17-10 Configuring Protocols for Deep Inspection in Inspection Rules (continued)

Protocol	Device Types	Policy Map	Class Map (ASA, PIX, FWSM only)	Description and Match Criteria Reference
H.323 H.225 H.323 RAS	ASA, PIX, FWSM	H.323 (ASA, PIX, FWSM)	H.323 (ASA, PIX, FWSM)	Inspect traffic based on a wide variety of criteria, including the H.323 message type, calling party, and called party. See Configuring H.323 Maps, page 17-51 and H.323 Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-54.
HTTP	ASA, PIX, FWSM, IOS	HTTP (ASA 7.1.x, PIX 7.1.x, FWSM 3.x, IOS) HTTP (ASA 7.2+, PIX 7.2+)	HTTP (ASA, PIX, FWSM)	Inspect traffic based on a wide variety of criteria including the content of the header or body, port misuse, and whether the traffic includes a Java applet. The maps used differ based on the operating system and version. **For ASA/PIX 7.2+**, see Configuring HTTP Maps for ASA 7.2+ and PIX 7.2+ Devices, page 17-64 and HTTP Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-66. **For ASA/PIX 7.1.x, FWSM 3.x+, and IOS**, see Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.
SIP	ASA, PIX, FWSM	SIP (ASA, PIX, FWSM)	SIP (ASA, PIX, FWSM)	Inspect traffic based on a wide variety of criteria. See Configuring SIP Maps, page 17-83 and SIP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-85.
Skinny	ASA, PIX, FWSM, IOS	Skinny	(none)	Inspect traffic based on a wide variety of criteria. See Configuring Skinny Maps, page 17-87 and Skinny Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-89.

Table 17-10 Configuring Protocols for Deep Inspection in Inspection Rules (continued)

Protocol	Device Types	Policy Map	Class Map (ASA, PIX, FWSM only)	Description and Match Criteria Reference
SMTP	ASA, PIX 7.x+, FWSM 3.x+, IOS	(none)	(none)	Inspect Simple Mail Transfer Protocol (SMTP) traffic and drop any that use illegal commands. You can configure a maximum data length for packets. See Configure SMTP Dialog Box, page 17-20.
SNMP	ASA, PIX, FWSM 3.x+, IOS	SNMP	(none)	Inspect SNMP traffic based on SNMP version. See Configuring SNMP Maps, page 17-90.
NetBIOS	ASA, PIX 7.x+, FWSM	NetBIOS	(none)	Inspect NetBIOS traffic to translate IP addresses in the NetBIOS name service (NBNS) packets according to the security appliance NAT configuration. You can drop packets that violate the protocol. See Configuring NetBIOS Maps, page 17-81.
IPSec Pass Through	ASA, PIX 7.x+	IPsec Pass Through	(none)	Inspect IPSec traffic and control whether ESP or AH traffic is allowed. See Configuring IPsec Pass Through Maps, page 17-80.
DCE/RPC	ASA 7.2+, PIX 7.2+, FWSM 3.2+	DCE/RPC	(none)	Inspect traffic based on timeouts and enforcing the mapper service. See Configuring DCE/RPC Maps, page 17-29.
IP options	ASA 8.2(2)+	IP Options	(none)	Allow IP packets that have certain options configured in the Options section of the IP header. In routed mode, packets that contain the router-alert option are allowed. Otherwise, if any option is set, packets are dropped. IP options are unnecessary for most communication, but the NOP (no operation) option might be used for padding, so you might want to allow it. See Configuring IP Options Maps, page 17-75.

Table 17-10 Configuring Protocols for Deep Inspection in Inspection Rules (continued)

Protocol	Device Types	Policy Map	Class Map (ASA, PIX, FWSM only)	Description and Match Criteria Reference
IPv6	ASA 8.4(2)+	IPv6	(none)	Inspect IPv6 traffic based on the following types of extension headers found anywhere in an IPv6 packet: Hop-by-Hop Options, Routing (Type 0), Fragment, Destination Options, Authentication, and Encapsulating Security Payload. See Configuring IPv6 Maps, page 17-77 and IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-78.
ESMTP	ASA, PIX 7.x+, FWSM 3.x+, IOS	ESMTP	(none)	Inspect ESMTP traffic. For IOS, you can configure only maximum data length. For ASA, PIX, FWSM, you can inspect traffic based on a wide variety of criteria. See Configuring ESMTP Maps, page 17-39.
Fragment	IOS	(none)	(none)	Inspect traffic based on a maximum allowed number of unassembled packet fragments. See Configure Fragments Dialog Box, page 17-20.
IMAP (Internet Message Access Protocol) POP3 (Post Office Protocol 3)	IOS	(none)	(none)	Inspect traffic based on invalid commands or clear text logins. See Configure IMAP or POP3 Dialog Boxes, page 17-21.
RPC (Sun Remote Procedure Call)	FWSM 2.x, IOS	(none)	(none)	Inspect traffic based on the RPC protocol number. See Configure RPC Dialog Box, page 17-21.

Table 17-10 Configuring Protocols for Deep Inspection in Inspection Rules (continued)

Protocol	Device Types	Policy Map	Class Map (ASA, PIX, FWSM only)	Description and Match Criteria Reference
IM	ASA, PIX 7.x+, IOS	IM (ASA 7.2+, PIX 7.2+) IM (IOS)	IM (only for ASA, PIX)	Inspect traffic based on a wide variety of criteria. The allowed maps differ based on operating system version. **For ASA, PIX**, see Configuring IM Maps for ASA 7.2+, PIX 7.2+ Devices, page 17-70 and IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-71. **For IOS**, see Configuring IM Maps for IOS Devices, page 17-73.
SCTP	ASA 9.5(2)+	SCTP	(none)	Inspect traffic based on Payload PID (PPID). See Configuring SCTP Maps, page 17-91 and SCTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-92
Diameter	ASA 9.5(2)+	Diameter	Diameter	Inspect traffic based on application ID, command codes, and AVP. See Configuring Diameter Maps, page 17-93 and Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-95
LISP	ASA 9.5(2)+	LISP	None	Inspect traffic allowed Endpoint Identifiers access list and validation key. See Configuring LISP Maps, page 17-102
M3UA	ASA 9.6(2)+	M3UA	None	Drops and logs packets that do not meet M3UA protocol conformance. See Configuring M3UA Maps, page 17-103

Related Topics

- Selecting Which Protocols To Inspect, page 17-3
- Understanding Inspection Rules, page 17-2
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5
- Configuring Inspection Rules, page 17-5
- Creating Policy Objects, page 6-9

- Understanding Map Objects, page 6-78
- Add/Edit Regular Expressions, page 17-108
- Configuring Regular Expression Groups, page 17-108

Configuring Class Maps for Inspection Policies

Use the Add and Edit Class Map dialog boxes to define class maps to be used in policy maps of the same type. The name of the dialog box indicates the type of map you are creating.

A class map defines application traffic based on criteria specific to the application. You then select the class map in the corresponding policy map and configure the action to take for the selected traffic. Thus, each class map must contain traffic that you want to handle in the same way (for example, to allow it or to drop it).

When configuring inspection rules for devices running ASA/PIX 7.2 or higher, or FWSM, you can create class maps for the inspection of the following types of traffic: DNS, FTP, H.323, HTTP, IM, SIP, and ScanSafe.

You can also define class criteria in the related policy map. However, creating class maps allows you to reuse the map in multiple policy maps.

The following topics describe the available match criteria:

- DNS Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-36
- FTP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-43
- H.323 Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-54
- HTTP Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-66
- IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-71
- SIP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-85
- Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-95

Navigation Path

Select **Manage > Policy Objects**, then select DNS, FTP, H.323 (ASA/PIX/FWSM), HTTP (ASA/PIX/FWSM), IM, SIP (ASA/PIX/FWSM), Diameter in the **Maps > Class Maps > Inspect** folder in the table of contents. Right-click inside the work area, then select **New Object**, or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Understanding Inspection Rules, page 17-2

Field Reference

Table 17-11 Add or Edit Class Maps Dialog Boxes for Inspection Rules

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Match table Match Type	The Match table lists the criteria included in the class map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion and the criterion and value that is inspected. • To add a criterion, click the Add button and fill in the Match Criterion dialog box. For more information, see the topics referenced above. • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring DCE/RPC Maps

Use the Add or Edit DCE/RPC Map dialog boxes to define a map for DCE/RPC inspection. A DCE/RPC inspection policy map lets you change the default configuration values used for DCE/RPC inspection.

DCE/RPC is a protocol widely used by Microsoft distributed client and server applications that allows software clients to execute programs on a server remotely.

This typically involves a client querying a server called the Endpoint Mapper listening on a well-known port number for the dynamically allocated network information of a required service. The client then sets up a secondary connection to the server instance providing the service. The security appliance allows the appropriate port number and network address and also applies NAT, if needed, for the secondary connection.

DCE/RPC inspection maps inspect for native TCP communication between the EPM and client on well-known TCP port 135. Map and lookup operations of the EPM are supported for clients. Client and server can be located in any security zone. The embedded server IP address and port number are received from the applicable EPM response messages. Because a client may attempt multiple connections to the server port returned by EPM, multiple use of pinholes are allowed, which have user configurable timeouts.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > DCE/RPC** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-12 Add and Edit DCE/RPC Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Pinhole Timeout	The timeout for DCE/RPC pinholes. The default is 2 minutes (00:02:00). Valid values are between 00:00:01 and 1193:00:00.
Enforce Endpoint Mapper Service	Whether to enforce the endpoint mapper service during binding. Using this service, a client queries a server, called the Endpoint Mapper, for the dynamically allocated network information of a required service.
Enable Endpoint Mapper Service Lookup Service Lookup Timeout	Whether to enable the lookup operation of the endpoint mapper service. If you select this option, you can enter the time out for the lookup operation. If you do not specify a timeout, the pinhole timeout or default pinhole timeout value is used. Valid values are between 00:00:01 and 1193:00:00.
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see DCE/RPC Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-31). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

DCE/RPC Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit DCE/RPC Match Criterion (for DCE/RPC class maps) or Match Condition and Action (for DCE/RPC policy maps) dialog boxes to do the following:

- Define the match criterion and value for a DCE/RPC class map.
- Select a DCE/RPC class map when creating a DCE/RPC policy map.
- Define the match criterion, value, and action directly in a DCE/RPC policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating a DCE/RPC class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for DCE/RPC, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating a DNS policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit DNS Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring DCE/RPC Maps, page 17-29.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-13 DCE/RPC Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing DCE/RPC class map or define a new DCE/RPC class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing DCE/RPC class map policy object. Enter the name of the DNS class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of traffic to match: • ms-rpc-epm—Matches Microsoft RPC EPM messages. • ms-rpc-isystemactivator—Matches ISystemMapper messages. • ms-rpc-oxidresolver-—Matches OxidResolver messages.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.

Table 17-13 DCE/RPC Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria. • Reset—Drop the packet, close the connection, and send a TCP reset to the server or client. • Log—Send a system log message. You can use this option alone or with one of the other actions. • Reset and Log— Perform the reset and log actions.

Configuring DNS Maps

Use the Add and Edit DNS Map dialog boxes to define DNS Maps for inspection. A DNS map lets you change the default configuration values used for DNS application inspection.

DNS application inspection supports DNS message controls that provide protection against DNS spoofing and cache poisoning. You can configure rules for certain DNS types to be allowed, dropped, or logged, while others are blocked. For example, you can restrict zone transfer between servers.

The Recursion Desired and Recursion Available flags in the DNS header can be masked to protect a public server from attack if that server only supports a particular internal zone. In addition, DNS randomization can be enabled to avoid spoofing and cache poisoning of servers that either do not support randomization or that use a weak pseudo random number generator. Limiting the domain names that can be queried protects the public server further.

You can configure a DNS mismatch alert as notification if an excessive number of mismatching DNS responses are received, which could indicate a cache poisoning attack.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > DNS** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-14 Add and Edit DNS Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.

Table 17-14 Add and Edit DNS Map Dialog Boxes (continued)

Element	Description
Protocol Conformance Tab	
	Defines DNS security settings and actions. For a description of the options on this tab, see DNS Map Protocol Conformance Tab, page 17-33.
Filtering Tab	
	Defines the filtering settings for DNS. For a description of the options on this tab, see DNS Map Filtering Tab, page 17-34.
Mismatch Rate Tab	
	The **Log When DNS ID Mismatch Rate Exceeds** option determines whether you want to report excessive instances of DNS identifier mismatches based on the following criteria: • Threshold—The maximum number of mismatch instances before a system message log is sent. Values are 0 to 4294967295. • Time Interval—The time period to monitor (in seconds). Values are 1 to 31536000.
Umbrella Connector Tab	
	Defines DNS umbrella connector settings for a DNS. For a description of the options on this tab, see DNS Umbrella Connector Tab, page 17-35.
Match Condition and Action Tab	
	The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see DNS Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-36). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

DNS Map Protocol Conformance Tab

Use the Protocol Conformance tab to define DNS security settings and actions for a DNS map.

Navigation Path

Click the Protocol Conformance tab on the Add and Edit DNS Map dialog boxes. See Configuring DNS Maps, page 17-32.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-15 DNS Map Protocol Conformance Tab

Element	Description
Enable DNS Guard Function	Whether to perform a DNS query and response mismatch check using the identification field in the DNS header. One response per query is allowed to go through the security appliance.
Generate Syslog for ID Mismatch	Whether to create syslog entries for excessive instances of DNS identifier mismatches.
Randomize the DNS Identifier for DNS Query	Whether to randomize the DNS identifier in the DNS query message.
Enable NAT Rewrite Function	Whether to enable IP address translation in the A record of the DNS response.
Enable Protocol Enforcement	Whether to enable DNS message format check, including domain name, label length, compression, and looped pointer check.
Enable DNS on TCP	Whether to enable inspection of DNS over TCP traffic.Ensure that DNS/TCP port 53 traffic is part of the class to which you apply DNS inspection. The inspection default class includes TCP/53.
Require Authentication Between DNS Server (RFC2845) Action	Whether to require authentication between DNS servers as defined in RFC 2845. If you select this option, select the action to take when there is no authentication.

DNS Map Filtering Tab

Use the Filtering tab to define DNS filtering settings and actions for a DNS map.

Navigation Path

Click the Filtering tab on the Add and Edit DNS Map dialog boxes. See Configuring DNS Maps, page 17-32.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-16 DNS Map Filtering Tab

Element	Description
Drop Packets that Exceed Specified Length Maximum Packet Length	Whether to drop packets that exceed the maximum length in bytes that you specify. This is a global setting.
Drop Packets Sent to Server that Exceed Specified Maximum Length Maximum Length	Whether to drop packets sent to the server that exceed the maximum length in bytes that you specify.
Drop Packets Sent to Server that Exceed Length Indicated by Resource Record	Whether to drop packets sent to the server that exceed the length indicated by the resource record.
Drop Packets Sent to Client that Exceed Specified Length Maximum Length	Whether to drop packets sent to a client that exceed the maximum length in bytes that you specify.
Drop Packets Sent to Client that Exceed Length Indicated by Resource Record	Whether to drop packets sent to the client that exceed the length indicated by the resource record.

DNS Umbrella Connector Tab

Use the Umbrella Connector tab to define DNS umbrella connector settings for a DNS map. Beginning with Cisco Security Manager version 4.18, the Umbrella global policy is supported on ASA 9.10.1 devices and above.

Navigation Path

Click the Umbrella Connector tab on the Add and Edit DNS Map dialog boxes. See Configuring DNS Maps, page 17-32.

Related Topics

- Configuring Umbrella Global Policy, page 48-16

Field Reference

Table 17-17 DNS Umbrella Connector Tab

Element	Description
Enable Umbrella Connector Tag for Umbrella Policy	Select the check box and enter the DNS policy-map umbrella tag name. The tag name can be a maximum of 50 characters. Cisco Security manager throws an error message if the tag name is greater than 50 characters. **Note** If the Umbrella global policy is not configured, Cisco Security Manager displays activity validation error. For more information on Umbrella global policy configuration, see Configuring Umbrella Global Policy, page 48-16.

Table 17-17 DNS Umbrella Connector Tab (continued)

Element	Description
Enable DNScrypt	Select this check box to enable the DNS crypt in the Umbrella datapath. For every hour, the secret key is exchanged between the key exchange thread and the Umbrella resolver. Ensure that the **Enable Umbrella Connector** check box is selected. If the check box is not selected, an error message is displayed for configuration discrepancy. **Note** If the Umbrella global policy is not configured, Cisco Security Manager displays activity validation error. For more information on Umbrella global policy configuration, see Configuring Umbrella Global Policy, page 48-16.

DNS Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit DNS Match Criterion (for DNS class maps) or Match Condition and Action (for DNS policy maps) dialog boxes to do the following:

- Define the match criterion and value for a DNS class map.
- Select a DNS class map when creating a DNS policy map.
- Define the match criterion, value, and action directly in a DNS policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating a DNS class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for DNS, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating a DNS policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit DNS Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring DNS Maps, page 17-32.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-18 DNS Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing DNS class map or define a new DNS class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing DNS class map policy object. Enter the name of the DNS class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of traffic to match: • DNS Class—Matches a DNS query or resource record class. • DNS Type—Matches a DNS query or resource record type. • Domain Name—Matches a domain name from a DNS query or resource record. • Header Flag—Matches a DNS flag in the header. • Question—Matches a DNS question. • Resource Record—Matches a DNS resource record.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Value (for DNS Class criterion)	The DNS class you want to inspect: • Internet—Matches the Internet DNS class. • DNS Class Field Value—Matches the specified number. • DNS Class Field Range—Matches the specified range of numbers.

Table 17-18 DNS Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Value (for DNS Type criterion)	The DNS type you want to inspect: • DNS Type Field Name—Matches the name of a DNS type: – A—IPv4 address. – AXFR—Full (zone) transfer. – CNAME—Canonical name. – IXFR—Incremental (zone) transfer. – NS—Authoritative name server. – SOA—Start of a zone of authority. – TSIG—Transaction signature. • DNS Type Field Value—Matches the specified number. • DNS Type Field Range—Matches the specified range of numbers.
Value (for Domain Name criterion)	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
Options Value (for Header Flag criterion)	The header flag you want to inspect. Use the Options field to indicate whether you want an exact match (Equals) or a partial match (Contains). • Header Flag Name—Matches the selected header flag names: – AA (authoritative answer) – QR (query) – RA (recursion available) – RD (recursion denied) – TC (truncation) flag bits • Header Flag Value—Matches the specified 16-bit hexadecimal value.
Resource Record	Lists the sections to match: • Additional—DNS additional resource record. • Answer—DNS answer resource record. • Authority—DNS authority resource record.

Configuring ESMTP Maps

Use the Add and Edit ESMTP Map dialog boxes to define the match criterion and values for the ESMTP inspect map. An ESMTP policy map lets you change the default configuration values used for ESMTP inspection.

ESMTP inspection detects attacks, including spam, phising, malformed message attacks, and buffer overflow/underflow attacks. It also provides support for application security and protocol conformance, which enforce the sanity of the ESMTP messages as well as detect several attacks, block senders/receivers, and block mail relay.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > ESMTP** from the Object Type selector. Right-click inside the table, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-19 Add and Edit ESMTP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Mask Server Banner	Whether to mask the server banner to prevent the client from discovering server information.
Configure Mail Relay Domain Name Action	Whether to have ESMTP inspection detect mail relay. When you select this option, enter the domain name you are inspecting and select the action you want to take when mail relay is detected.
Special Character (ASA7.2.3+/PIX7.2.3+) Action	Whether you want to detect special characters in sender or receiver email addresses. If you select this option, select the action you want to take when special characters are detected.
Allow TLS (ASA7.2.3+, 8.0.3+/PIX7.2.3) Action Log	Whether to allow a TLS proxy on the security appliance. If you select this option, you can also select Action Log to create a log entry when TLS is detected.

Table 17-19 Add and Edit ESMTP Map Dialog Boxes (continued)

Element	Description
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see ESMTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-40). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

ESMTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Add or Edit Match Condition and Action dialog boxes to define the match criterion, value, and action for an ESMTP policy map.

The fields on this dialog box change based on the criterion you select. You can use the following criteria:

- Body Length—Matches the message body length.
- Body Line Length—Matches the length of a line in the message body.
- Commands—Matches ESMTP commands.
- Command Recipient Count—Matches the number of recipient email addresses.
- Command Line Length—Matches the number of characters of a command line.
- EHLO Reply Parameters—Matches the ESMTP EHLO reply parameters.
- Header Length—Matches the number of characters of the header.
- Header Line Length—Matches the number of characters of a line in the message header.
- To Recipients Count—Matches the number of recipients in the To field of the header.
- Invalid Recipients Count—Matches the number of invalid recipients in the header.
- MIME File Type—Matches the MIME file type.
- MIME Filename Length—Matches the number of characters of the filename.
- MIME Encoding—Matches the MIME encoding scheme.
- Sender Address—Matches the address of the sender.
- Sender Address Length—Matches the number of characters of the sender's address.

Navigation Path

In the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit ESMTP Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring ESMTP Maps, page 17-39.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-20 ESMTP Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Specifies which criterion of ESMTP traffic to match. The criteria are described above.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action	The action you want the device to take for traffic that matches the defined criteria.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Greater Than Length	The length in bytes of the evaluated field. The criterion matches if the length is greater than the specified number, and does not match if the field is less than the specified number. The dialog box indicates the valid range for the length, except for Body Length and Header length, which can be 1 to 4294967295.
Commands	The ESMTP command verbs you want to inspect.
Greater Than Count	The number of evaluated items. The criterion matches if the count is greater than the specified number, and does not match if the count is less than the specified number.
Parameters	The ESMTP EHLO reply parameters you want to inspect.

Table 17-20 ESMTP Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
MIME Encoding	The type of MIME encoding schemes you want to inspect.

Configuring FTP Maps

Use the Add and Edit FTP Map dialog boxes to define the match criterion and values for an FTP inspect map. You can use an FTP map to block specific FTP protocol methods, such as an FTP PUT, from passing through the security appliance and reaching your FTP server.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > FTP** from the Object Type selector. Right-click inside the table, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-21 Add and Edit FTP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Mask Greeting Banner from Server	Whether to mask the greeting banner from the FTP server to prevent the client from discovering server information.
Mask Reply to SYST Command	Whether to mask the reply to the **syst** command to prevent the client from discovering server information.

Table 17-21 Add and Edit FTP Map Dialog Boxes (continued)

Element	Description
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see FTP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-43). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.
Validate For Validate button	The device platforms for which to validate the object. Select the platform for which you intend to use this object and click **Validate** to determine if the object is configured in a way that will prevent policy deployment.

FTP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit FTP Match Criterion (for FTP class maps) or Match Condition and Action (for FTP policy maps) dialog boxes to do the following:

- Define the match criterion and value for an FTP class map.
- Select an FTP class map when creating an FTP policy map.
- Define the match criterion, value, and action directly in an FTP policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating an FTP class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for FTP, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating an FTP policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit FTP Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring FTP Maps, page 17-42.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-22 FTP Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing FTP class map or define a new FTP class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing FTP class map policy object. Enter the name of the FTP class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of FTP traffic to match: • Request Command—Matches an FTP request command. • Filename—Matches a filename for FTP transfer. • File Type—Matches a file type for FTP transfer. • Server—Matches an FTP server name. • Username—Matches an FTP username.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.

Variable Fields

The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.

Table 17-22 FTP Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Request Commands	The FTP commands you want to inspect: • Append (APPE)—Appends to a file. • Delete (DELE)—Deletes a file at the server site. • Help (HELP)—Provides help information from the server. • Put (PUT)—FTP client command for the stor (store a file) command. • Rename From (RNFR)—Specifies rename-from filename. • Server Specific Command (SITE)—Specifies commands that are server specific. Usually used for remote administration. • Change to Parent (CDUP)—Changes to the parent directory of the current working directory. • Get (GET)—FTP client command for the retr (retrieve a file) command. • Create Directory (MKD)—Creates a directory. • Remove Directory (RMD)—Removes a directory. • Rename To (RNTO)—Specifies rename-to filename. • Store File with Unique Name (STOU)—Stores a file with a unique filename.
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.

Configuring GTP Maps

Use the Add and Edit GTP Map dialog boxes to define the match criterion and values for a GTP inspect map.

The GPRS Tunnel Protocol (GTP) provides uninterrupted connectivity for mobile subscribers between GSM networks and corporate networks or the Internet. GTP uses a tunneling mechanism to provide a service for carrying user data packets.

A GTP map object lets you change the default configuration values used for GTP application inspection. The GTP protocol is designed to provide security for wireless connections to TCP/IP networks such as the Internet. You can use a GTP map to control timeout values, message sizes, tunnel counts, and GTP versions traversing the security appliance.

Starting from version 4.18, Cisco Security Manager supports anti-replay feature of ASA 9.10.1. By enabling data packet replay, your network is protected from replay attacks.

Tip GTP inspection requires a special license. If you do not have the required license, you will see device errors if you try to deploy a GTP map.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > GTP** from the Object Type selector. Right-click inside the work area, then select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-23 Add and Edit GTP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Country and Network Codes Table	The three-digit Mobile Country Code (mcc) and Mobile Network Code (mnc) to include in the map. The codes are 000 to 999. • To add codes, click the Add button and fill in the dialog box. • To edit a row, select it and click the Edit button. • To delete a row, select it and click the Delete button.
Permit Response Table	The Network/Host policy objects for which you will allow GTP responses from a GSN that is different from the one to which the response was sent. • To add objects, click the Add button and fill in the dialog box. For more information, see Add and Edit Permit Response Dialog Boxes, page 17-48. • To edit a row, select it and click the Edit button. • To delete a row, select it and click the Delete button.
Request Queue	The maximum requests allowed in the queue. When the limit has been reached and a new request arrives, the request that has been in the queue for the longest time is removed. Values are 1-9999999. The default is 200.

Table 17-23 Add and Edit GTP Map Dialog Boxes (continued)

Element	Description
Tunnel Limit	The maximum number of tunnels allowed.
Permit Errors	Whether to permit packets with errors or different GTP versions. By default, all invalid packets or packets that failed during parsing are dropped.
Enable Data Packet Replay Window	Select the check box to configure the anti-replay and select one of the four window sizes—128, 256, 512, or 1024. Messages that are outside of the window size are dropped. For information on configuration of GTP Map policy, refer to Add or Edit Inspect/Application FW Rule Wizard, Inspected Protocol Page, page 17-17.
Enable Header	Check Select the check box to enable header check of the data packets.
Anti-User Spoofing	This field is enabled only when you select the Enable Header Check check box. Select the relevant option: • Bypass—to forward the packets that pass the header check. • Drop—to drop the packets that pass the header check.
Edit Timeouts button	Click this button to configure time out values for various operations. For more information about the options, see GTP Map Timeouts Dialog Box, page 17-48.

Match Condition and Action Tab

The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions.

- To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see GTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-49).
- To edit a criterion, select it and click the Edit button.
- To delete a criterion, select it and click the Delete button.

Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.
Validate For Validate button	The device platforms for which to validate the object. Select the platform for which you intend to use this object and click **Validate** to determine if the object is configured in a way that will prevent policy deployment.

Add and Edit Country Network Codes Dialog Boxes

Use the Add and Edit Country Network Codes dialog boxes to add Mobile Country Code (mcc) and Mobile Network Code (mnc) values to the GTP policy map. The codes can be 000 to 999.

Navigation Path

From the Add and Edit GTP Map dialog boxes, click the **Add** button in the Country and Network codes table, or select a row and click the **Edit** button. See Configuring GTP Maps, page 17-45.

Add and Edit Permit Response Dialog Boxes

Use the Add and Edit Permit Response dialog boxes to permit GTP responses from a GSN that is different from the one to which the response was sent.

Enter the name of a Network/Host policy object that defines the destination (**To Object Group**) and source (**From Object Group**) of the traffic. You can click **Select** to select the object from a list, where you can also create an new object by clicking the **Create** button in the Object Selector dialog box.

You cannot use the Network/Host object named “any.”

Navigation Path

From the Add and Edit GTP Map dialog boxes, click the **Add** button in the Permit Response table, or select a row and click the **Edit** button. See Configuring GTP Maps, page 17-45.

GTP Map Timeouts Dialog Box

Use the GTP Map Timeouts dialog box to set timeout values for a GTP Map.

Navigation Path

From the Add and Edit GTP Map dialog boxes, click the **Edit Timeouts** button on the Parameters tab. See Configuring GTP Maps, page 17-45.

Field Reference

Table 17-24 GTP Map Timeouts Dialog Box

Element	Description
GSN Timeout (Prior to ASA 9.5(1)) Endpoint Timeout (ASA 9.5(1) or higher)	The period of inactivity (hh:mm:ss) after which a GSN is removed. The default is 30 minutes. Enter 0 to never tear down immediately.
PDP Context Timeout	The maximum period of time allowed (hh:mm:ss) before beginning to receive the PDP context. The default is 30 minutes. Enter 0 to specify no limit.
Request Queue Timeout	The maximum period of time allowed (hh:mm:ss) before beginning to receive the GTP message. The default is 60 seconds. Enter 0 to specify no limit.
Signaling Connections Timeout	The period of inactivity (hh:mm:ss) after which the GTP signaling is removed. The default is 30 minutes. Enter 0 to not remove the signal.

Table 17-24 GTP Map Timeouts Dialog Box (continued)

Element	Description
Tunnel Timeout	The period of inactivity (hh:mm:ss) after which the GTP tunnel is torn down. The default is 60 seconds (when a Delete PDP Context Request is not received). Enter 0 to never tear down immediately.
T3 Response Timeout	The maximum wait time for a response before removing the connection.

GTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Add or Edit Match Condition and Action dialog boxes to define the match criterion, value, and action for a GTP policy map.

The fields on this dialog box change based on the criterion you select.

Navigation Path

In the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit GTP Map dialog box, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring GTP Maps, page 17-45.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-25 GTP Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Specifies which criterion of GTP traffic to match: • Access Point Name—Matches the access point name so you can define the access points to drop when GTP application inspection is enabled. • Message ID—Matches the numeric identifier for the message that you want to drop. By default, all valid message IDs are allowed. • Message Length—Matches the length of the UDP packet. Use this criterion to change the default for the maximum allowed message length for the UDP payload. • Version—Matches the GTP version. • MSISDN—Matches the MSISDN with regular expressions or class and drop all GTP packets that have matching MSISDN. • Selection Mode—Ranges between 0 and 3.

Table 17-25 GTP Policy Maps Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action	The action you want the device to take for traffic that matches the defined criteria. • Drop Packet—By default, all invalid packets or packets that failed during parsing are dropped. • Drop Packet and Log • Rate Limit
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Access Point Name	The access points to act on when GTP application inspection is enabled. • Specified By—An access point name to be dropped. By default, all messages with valid APNs are inspected, and any APN is allowed. • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
ID type	The numeric identifier of the message that you want to act on. • Value—A single message ID. The Value can be between 1 and 255. • Range—A range of message IDs. The Range can be between 1 and 255.
Minimum Length	The minimum number of bytes in the UDP payload.
Maximum Length	The maximum number of bytes in the UDP payload.

Table 17-25 GTP Policy Maps Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Version	Beginning with version 4.9, Security Manager provides support for GPRS Tunnel Protocol (GTP) v2 and enhanced v1 in the GTP Map Object for ASA devices 9.5(1) or higher. You can now configure separate message ID matching for GTPv1 and GTPv2. For ASA devices 9.5(1) or higher, if you select Message ID as the Criterion, two options for Version, v1 and v2, are displayed. Select v1 or v2 and enter a single Value between 1 and 255, or a Range of values from 1 to 255.
Version Type	Prior to ASA version 9.5(1)—Use 0 to identify Version 0 and 1 to identify Version 1. Version 0 of GTP uses port 2123, while Version 1 uses port 3386. By default all GTP versions are allowed.
Regular Expression	Beginning with version 4.18, Cisco Security Manager allows configuring of MSISDN with regular expressions and drop all GTP packets that have matching MSISDN. This field appears when you select MSISDN in the Criterion drop-down.
Regular Expression Group	Beginning with version 4.18, Cisco Security Manager allows configuring of MSISDN with regular expressions class and drop all GTP packets that have matching MSISDN. This field appears when you select MSISDN in the Criterion drop-down.
Mode Value	If Selection is selected in the Criterion drop-down, this field appears. Enter the mode value in the range of 0 – 3. This is a mandatory field.

Configuring H.323 Maps

Use the Add and Edit H.323 Map dialog boxes to define the match criterion and values for an H.323 inspect map. An H.323 policy map lets you change the default configuration values used for H.323 inspection.

H.323 inspection supports H.323 compliant applications such as Cisco CallManager and VocalTec Gatekeeper. H.323 is a suite of protocols defined by the International Telecommunication Union for multimedia conferences over LANs. The security appliance supports H.323 through Version 4, including H.323 v3 feature Multiple Calls on One Call Signaling Channel.

With H.323 inspection enabled, the security appliance supports multiple calls on the same call signaling channel, a feature introduced with H.323 Version 3. This feature reduces call setup time and reduces the use of ports on the security appliance. The two major functions of H.323 inspection are as follows:

- NAT the necessary embedded IPv4 addresses in the H.225 and H.245 messages. Because H.323 messages are encoded in PER encoding format, the security appliance uses an ASN.1 decoder to decode the H.323 messages.
- Dynamically allocate the negotiated H.245 and RTP/RTCP connections.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > H.323 (ASA/PIX/FWSM)** from the Object Type selector. Right-click inside the work area, then select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-26 Add and Edit H.323 Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
HSI Group table	The HSI groups to include in the map. The group number, IP address of the HSI host, and IP addresses and interface names of the clients connected to the security appliance are shown in the table. Up to five HSI hosts per group, and up to ten end points per HSI group, are allowed. • To add a group, click the Add button and fill in the dialog box (see Add or Edit HSI Group Dialog Boxes, page 17-53). • To edit a group, select it and click the Edit button. • To delete a group, select it and click the Delete button.
Call Duration Limit	The call duration limit in seconds. The range is from 0:0:0 to 1163:0:0. A value of 0 means never timeout.
Enforce Presence of Calling and Called Party Numbers	Whether to enforce calling and called party numbers used in call setup.
Allow the facility message before SETUP for H.460.18	Whether to allow the FACILITY message to be sent before the SETUP message as part of the Incoming Call Message Procedure. **Note** H.460.18 defines a method for traversal of H.323 signaling across network address translators and firewalls.
Check State Transition on H.225 Messages	Whether to enable state checking validation on H.225 messages.
Check State Transition on RAS Messages	Whether to enable state checking validation on RAS messages.
Create Pinholes on Seeing RCF Packets	Whether to enable call setup between H.323 endpoints when the Gatekeeper is inside the network. The device opens pinholes for calls based on Registration Request/Registration Confirm (RRQ/RCF) messages. Because these RRQ/RCF messages are sent to and from the Gatekeeper, the calling endpoint's IP address is unknown and the device opens a pinhole through source IP address/port 0/0. This option is available for ASA 8.0(5)+ devices.
Check for H.245 Tunneling Action	Whether to enforce H.245 tunnel blocking and perform the action you select in the Action list box.

Table 17-26 Add and Edit H.323 Map Dialog Boxes (continued)

Element	Description
Check RTP Packets for Protocol Conformance	Whether to check RTP packets flowing through the pinholes for protocol conformance.
Payload Type must be Audio or Video based on Signaling Exchange	Whether to enforce the payload type to be audio or video based on the signaling exchange.

Match Condition and Action Tab

The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions.

- To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see H.323 Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-54).
- To edit a criterion, select it and click the Edit button.
- To delete a criterion, select it and click the Delete button.

Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit HSI Group Dialog Boxes

Use the Add or Edit HSI group dialog boxes to add HSI groups to an H.323 policy inspection map.

Navigation Path

From the Parameters tab on the Add and Edit H.323 Map dialog boxes, click the **Add Row** button in the HSI group table, or select a row and click the **Edit Row** button. See Configuring H.323 Maps, page 17-51.

Field Reference

Table 17-27 Add and Edit HSI Group Dialog Boxes

Element	Description
Group ID	The HSI group ID number (0 to 2147483647).
IP Address	The IP address of the HSI host.

Table 17-27 Add and Edit HSI Group Dialog Boxes (continued)

Element	Description
Endpoint table	The end points associated with HSI group. You can add up to 10 end points per group. For each end point, you specify the IP address and interface policy group. • To add an end point, click the Add button and fill in the dialog box (see Add or Edit HSI Endpoint IP Address Dialog Boxes, page 17-54). • To edit an end point, select it and click the Edit button. • To delete an end point, select it and click the Delete button.

Add or Edit HSI Endpoint IP Address Dialog Boxes

Us the Add or Edit HSI Endpoint IP Address dialog box to add end points to an HSI group.

Navigation Path

From the Add and Edit HSI Group dialog boxes, click the **Add Row** button in the end point table, or select a row and click the **Edit Row** button. See Configuring H.323 Maps, page 17-51.

Field Reference

Table 17-28 Add and Edit HSI Endpoint IP Address Dialog Boxes

Element	Description
Network/Host	The IP address of the end point host or network.
Interface	The Interface policy group that identifies the interface connected to the security appliance. Enter the name of a policy group, or click **Select** to select it from a list, where you can also create new policy groups.

H.323 Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit H.323 Match Criterion (for H.323 class maps) or Match Condition and Action (for H.323 policy maps) dialog boxes to do the following:

- Define the match criterion and value for an H.323 class map.
- Select an H.323 class map when creating an H.323 policy map.
- Define the match criterion, value, and action directly in an H.323 policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating an H.323 class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for H.323, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating an H.323 policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit H.323 Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring H.323 Maps, page 17-51.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-29 H.323 Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing H.323 class map or define a new H.323 class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing H.323 class map policy object. Enter the name of the H.323 class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of H.323 traffic to match: • Called Party—Matches the called party address. • Calling Party—Matches the calling party address. • Media Type—Matches the media type.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
Media Type	The type of media you want to inspect, audio, video, or data.

Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX, FWSM, and IPS features/functionality, it does not support any enhancements.

Use the Add and Edit HTTP Map dialog boxes to define HTTP maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x, and IOS devices.

The enhanced HTTP inspection feature, which is also known as an application firewall, verifies that HTTP messages conform to RFC 2616, use RFC-defined methods, and comply with various other criteria. This can help prevent attackers from using HTTP messages for circumventing network security policy.

When you enable HTTP inspection with an HTTP map, strict HTTP inspection with the action reset and log is enabled by default. You can change the actions performed in response to inspection failure, but you cannot disable strict inspection as long as the HTTP map remains enabled. Security Manager uses the **http-map** command to configure the map on the device.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > HTTP (ASA 7.1.x/PIX 7.1.x/FWSM3.x/IOS)** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-30 Add and Edit HTTP Map Dialog Boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
General tab	Defines the action taken when non-compliant HTTP requests are received and to enable verification of content type. For a description of the options, see HTTP Map General Tab, page 17-57.
Entity Length tab	Defines the action taken if the length of the HTTP content falls outside of configured targets. For a description of the options, see HTTP Map Entity Length Tab, page 17-58.
RFC Request Method tab	Defines the action that the security appliance should take when specific RFC request methods are used in the HTTP request. For a description of the options, see HTTP Map RFC Request Method Tab, page 17-60.
Extension Request Method tab	Defines the action taken when specific extension request methods are used in the HTTP request. For a description of the options, see HTTP Map Extension Request Method Tab, page 17-61.

Table 17-30 Add and Edit HTTP Map Dialog Boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices

Element	Description
Port Misuse tab	Defines the action taken when specific undesirable applications are encountered. For a description of the options, see HTTP Map Port Misuse Tab, page 17-62.
Transfer Encoding tab	Defines the action taken when specific transfer encoding types are used in the HTTP request. For a description of the options, see HTTP Map Transfer Encoding Tab, page 17-63.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

HTTP Map General Tab

Use the General tab to define the action taken when non-compliant HTTP requests are received and to enable verification of content type.

Navigation Path

Click the General tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-31 HTTP Map General Tab

Element	Description
Take action for non-RFC 2616 compliant traffic	Whether you want to configure the action to be taken for traffic that does not comply with RFC 2616. Possible actions are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server. You can also select **Generate Syslog** to write a message to the syslog if non-compliant traffic is encountered.

Table 17-31 HTTP Map General Tab (continued)

Element	Description
Verify Content-type field belongs to the supported internal content-type list.	Whether you want to configure the action to be taken for traffic whose content type does not belong to the supported internal content-type list. Possible actions are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server. You can also select these options: • **Verify Content-type field for response matches the ACCEPT field of request**—To also verify that the content type of the response matches the request. • **Generate Syslog**—To write a message to the syslog if non-compliant traffic is encountered.
Override Global TCP Idle Timeout (IOS only)	Whether to change the TCP idle timeout default setting. An IOS device terminates a connection if there is no communication activity after this length of time. If you select this option, specify the desired timeout value in seconds.
Override Global Audit Trail Setting (IOS only) Enable Audit Trail	Whether to change the audit trail setting for IOS devices. If you select this option, you can select **Enable Audit Trail** to generate audit trail messages.

HTTP Map Entity Length Tab

Use the Entity Length tab to enable inspection based on the length of the HTTP content.

Navigation Path

Click the Entity Length tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-32 HTTP Map Entity Length Tab

Element	Description
Inspect URI Length	Whether to enable inspection based on the length of the URI. If you select this option, configure the following: • Maximum—The desired maximum length, in bytes, of the URI, from 1 to 65535. • Excessive URI Length Action—The action to take when the length is exceeded: – Allow Packet—Allow the message. – Drop Packet—Close the connection. – Reset Connection—Send a TCP reset message to client and server. • Generate Syslog—Whether to generate a syslog message when a violation occurs.
Inspect Maximum Header Length	Whether to enable inspection based on the length of the HTTP header. If you select this option, configure the following: • Request—The desired maximum length, in bytes, of the request header, from 1 to 65535. • Response—The desired maximum length, in bytes, of the response header, from 1 to 65535. • Excessive Header Length Action—The action to take when the length is exceeded: – Allow Packet—Allow the message. – Drop Packet—Close the connection. – Reset Connection—Send a TCP reset message to client and server. • Generate Syslog—Whether to generate a syslog message when a violation occurs.

Table 17-32 HTTP Map Entity Length Tab (continued)

Element	Description
Inspect Body Length	Whether to enable inspection based on the length of the message body. If you select this option, configure the following: • Minimum Threshold—The desired minimum length, in bytes, of the message body, from 1 to 65535. • Maximum Threshold—The desired maximum length, in bytes, of the message body, from 1 to 65535. • Body Length Threshold Action—The action to take when the message body falls outside of the configured boundaries: – Allow Packet—Allow the message. – Drop Packet—Close the connection. – Reset Connection—Send a TCP reset message to client and server. • Generate Syslog—Whether to generate a syslog message when a violation occurs.

HTTP Map RFC Request Method Tab

Use the RFC Request Method tab to define the action to take when specific request methods are used in the HTTP request.

Navigation Path

Click the RFC Request Method tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-33 HTTP Map RFC Request Method

Element	Description
Available and Selected Methods	The Available Methods list contains the request methods defined in RFC 2616.
Action Generate Syslog	To configure an action for a method, select it, then select an action and optionally select **Generate Syslog** if you want a message added to the syslog when an HTTP request containing the selected method is encountered. Click the **>>** button to add it to the Selected Methods list. (To remove a method from the selected list, select it and click the **<<** button.) **Tip** You can select multiple methods at a time using Ctrl+click if the action and syslog requests are the same for each. The actions you can specify are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server.
Specify the action to be applied for the remaining available methods above.	Whether to define a default action for the methods for which you have not configured specific actions above. If you select this option, select the action and syslog setting to use for the default action.

HTTP Map Extension Request Method Tab

Use the Extension Request Method tab to define the action taken when specific extension request methods are used in the HTTP request.

Navigation Path

Click the Extension Request Method tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-34 HTTP Map Extension Request Method Tab

Element	Description
Available and Selected Methods	The Available Methods list contains the extension request methods defined in RFC 2616.
Action Generate Syslog	To configure an action for a method, select it, then select an action and optionally select **Generate Syslog** if you want a message added to the syslog when an HTTP request containing the selected method is encountered. Click the >> button to add it to the Selected Methods list. (To remove a method from the selected list, select it and click the << button.) **Tip** You can select multiple methods at a time using Ctrl+click if the action and syslog requests are the same for each. The actions you can specify are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server.
Specify the action to be applied for the remaining available methods above.	Whether to define a default action for the methods for which you have not configured specific actions above. If you select this option, select the action and syslog setting to use for the default action.

HTTP Map Port Misuse Tab

Use the Port Misuse tab to enable port misuse application firewall inspection. The application categories you can configure are:

- IM—Instant Messaging. The applications checked for are Yahoo! Messenger, AIM, and MSN IM.
- P2P—Peer-to-peer applications. The Kazaa application is checked.
- Tunneling—Tunneling applications. The applications checked for are HTTPort/HTTHost, GNU Httptunnel, GotoMyPC, Firethru, and Http-tunnel.com Client.

Navigation Path

Click the Port Misuse tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-35 HTTP Map Port Misuse Tab

Element	Description
Available and Selected Application Categories Action Generate Syslog	The Available Application Categories list contains the categories for which you can define firewall inspection settings. To configure an action for a category, select it, then select an action and optionally select **Generate Syslog** if you want a message added to the syslog when an HTTP request containing the selected application is encountered. Click the >> button to add it to the Selected Categories list. (To remove a category from the selected list, select it and click the << button.) **Tip** You can select multiple categories at a time using Ctrl+click if the action and syslog requests are the same for each. The actions you can specify are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server.
Specify the action to be applied for the remaining available categories above.	Whether to define a default action for the categories for which you have not configured specific actions above. If you select this option, select the action and syslog setting to use for the default action.

HTTP Map Transfer Encoding Tab

Use the Transfer Encoding tab to enable inspection based on the transfer encoding type. The encoding types that you can configure are:

- Chunked—Identifies the transfer encoding type in which the message body is transferred as a series of chunks.
- Compressed—Identifies the transfer encoding type in which the message body is transferred using UNIX file compression.
- Deflate—Identifies the transfer encoding type in which the message body is transferred using zlib format (RFC 1950) and deflate compression (RFC 1951).
- GZIP—Identifies the transfer encoding type in which the message body is transferred using GNU zip (RFC 1952).
- Identity—Identifies connections in which no transfer encoding is performed in the message body.

Navigation Path

Click the Transfer Encoding tab on the Add and Edit HTTP Map dialog boxes for ASA 7.1.x/PIX 7.1.x/FWSM 3.x/IOS Devices. See Configuring HTTP Maps for ASA 7.1.x, PIX 7.1.x, FWSM 3.x and IOS Devices, page 17-56.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-36 HTTP Map Transfer Encoding Tab

Element	Description
Available and Selected Encoding Types	The Available Encoding Types list contains the types of transfer encoding for which you can define firewall inspection settings.
Action Generate Syslog	To configure an action for a type, select it, then select an action and optionally select **Generate Syslog** if you want a message added to the syslog when an HTTP request containing the selected type is encountered. Click the **>>** button to add it to the Selected Encoding Types list. (To remove a type from the selected list, select it and click the **<<** button.) **Tip** You can select multiple types at a time using Ctrl+click if the action and syslog requests are the same for each. The actions you can specify are: • Allow Packet—Allow the message. • Drop Packet—Close the connection. • Reset Connection (default)—Send a TCP reset message to client and server.
Specify the action to be applied for the remaining available encoding types above.	Whether to define a default action for the types for which you have not configured specific actions above. If you select this option, select the action and syslog setting to use for the default action.

Configuring HTTP Maps for ASA 7.2+ and PIX 7.2+ Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the Add and Edit HTTP Map dialog boxes to define the match criterion and values for the HTTP inspect map for ASA and PIX software releases 7.2 and higher.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > HTTP (ASA 7.2+/PIX 7.2+)** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-37 Add and Edit HTTP Map Dialog Boxes (ASA 7.2+/PIX 7.2+)

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Body Match Maximum	The maximum number of characters in the body of an HTTP message that should be searched in a body match. **Tip** A high value can have a significant impact on performance.
Check for protocol violations	Whether to check for protocol violations.
Action	The action to take based on the defined settings. You can drop, reset, or log the connection.
Spoof Server	Enables you to replace the server HTTP header value with the specified string.

Match Condition and Action Tab

The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions.

- To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see HTTP Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-66).
- To edit a criterion, select it and click the Edit button.
- To delete a criterion, select it and click the Delete button.

Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.
Overrides: None	Shows that no overrides exist on the device. You must manually set overrides in order to change the display. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18. **Note** Selecting Allow Value Override per Device does not automatically set overrides.

HTTP Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the Add or Edit HTTP Match Criterion (for HTTP class maps) or Match Condition and Action (for HTTP policy maps) dialog boxes to do the following:

- Define the match criterion and value for an HTTP class map.
- Select an HTTP class map when creating an HTTP policy map.
- Define the match criterion, value, and action directly in an HTTP policy map.

These types of maps are used only for devices running ASA 7.2 or higher, or PIX 7.2 or higher, operating systems.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map. You can use the following criteria:

- Request/Response Content Type Mismatch—Specifies that the content type in the response must match one of the MIME types in the accept field of the request.
- Request Arguments—Applies the regular expression match to the arguments of the request.
- Request Body—Applies the regular expression match to the body of the request.
- Request Body Length—Specifies that the body length of the request be matched as greater than or less than the specified number of bytes.
- Request Header Count—Specifies that the number of headers in the request be matched as greater than or less than the specified number.
- Request Header Length—Specifies that the header length of the request be matched as greater than or less than the specified number of bytes.
- Request Header Field—Applies the regular expression match to the header of the request.
- Request Header Field Count—Applies the regular expression match to the header of the request based on a specified number of header fields.
- Request Header Field Length—Applies the regular expression match to the header of the request based on a specified field length.
- Request Header Content Type—Specifies the content type to evaluate in the content-type header field of the request.
- Request Header Transfer Encoding—Specifies the transfer encoding to evaluate in the transfer-encoding header field of the request.
- Request Header Non-ASCII—Specifies whether there are non-ASCII characters in the header of the request.
- Request Method—Specifies the method of the request to match.
- Request URI—Applies the regular expression match to the URI of the request.
- Request URI Length—Specifies that the URI length of the request be matched as greater than or less than the specified number of bytes.
- Response Body ActiveX—Specifies whether there is ActiveX content in the body of the request.
- Response Body Java Applet—Specifies whether there is a Java applet in the body of the request.

- Response Body—Applies the regular expression match to the body of the response.
- Response Body Length—Specifies that the body length of the response be matched as greater than or less than the specified number of bytes.
- Response Header Count—Specifies that the number of headers in the response be matched as greater than or less than the specified number.
- Response Header Length—Specifies that the header length of the response be matched as greater than or less than the specified number of bytes.
- Response Header Field—Applies the regular expression match to the header of the response.
- Response Header Field Count—Applies the regular expression match to the header of the response based on a specified number of header fields.
- Response Header Field Length—Applies the regular expression match to the header of the response based on a specified field length.
- Response Header Content Type—Specifies the content type to evaluate in the content-type header field of the response.
- Response Header Transfer Encoding—Specifies the transfer encoding to evaluate in the transfer-encoding header field of the response.
- Response Header Non-ASCII—Specifies whether there are non-ASCII characters in the header of the response.
- Response Status Line—Applies the regular expression match to the status line of the response.

Navigation Path

When creating an HTTP class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for HTTP, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating an HTTP policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit HTTP Map dialog boxes for ASA/PIX 7.2+ devices, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring HTTP Maps for ASA 7.2+ and PIX 7.2+ Devices, page 17-64.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-38 HTTP Class and Policy Maps (ASA 7.2+/PIX 7.2+) Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing HTTP class map or define a new HTTP class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing HTTP class map policy object. Enter the name of the HTTP class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of HTTP traffic to match. The criteria are described above.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. For some criteria, this is the only available option. • Doesn't Match—Does not match the criterion.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria. The types of action depend on the criterion you select.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Field Name	The name of the header field to evaluate. You can select one of the following: • Predefined—The predefined HTTP header fields. • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object.

Table 17-38 HTTP Class and Policy Maps (ASA 7.2+/PIX 7.2+) Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object. When you are evaluating the Request Header Transfer Encoding or Response Header Transfer Encoding criteria, you can also select these options: • Specified By—One of the following predefined types of transfer encoding: – Chunked—The message body is transferred as a series of chunks. – Compressed—The message body is transferred using UNIX file compression. – Deflate—The message body is transferred using zlib format (RFC 1950) and deflate compression (RFC 1951). – GZIP—The message body is transferred using GNU zip (RFC 1952). – Identity—No transfer encoding is performed. • Empty—The transfer-encoding field in request header is empty.
Greater Than Length	The length in bytes of the evaluated field. The criterion matches if the length is greater than the specified number, and does not match if the field is less than the specified number.
Greater Than Count	The number of evaluated items. The criterion matches if the count is greater than the specified number, and does not match if the count is less than the specified number.

Table 17-38 HTTP Class and Policy Maps (ASA 7.2+/PIX 7.2+) Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Content Type	The content type to evaluate as specified in the content-type header field. You can select one of the following: • Specified By—A predefined MIME type. • Unknown—The MIME type is not known. Select Unknown when you want to evaluate the item against all known MIME types. • Violation—The magic number in the body must correspond to the MIME type in the content-type header field. • Regular Expression, Regular Expression Group—The regular expression or regular expression group to evaluate. See the explanation for the Value field for an explanation of these options.
Request Method	The specified request method to match. You can select one of the following: • Specified By—The predefined request method. • Regular Expression, Regular Expression Group—The regular expression or regular expression group to evaluate. See the explanation for the Value field for an explanation of these options.

Configuring IM Maps for ASA 7.2+, PIX 7.2+ Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the Add and Edit IM Map dialog boxes to define settings for define an Instant Messenger (IM) inspect map for devices running ASA/PIX 7.2 or higher. An IM map lets you change the default configuration values used for IM application inspection.

Instant Messaging causes concern due to its use of clear text when conducting business and the potential for network attacks and the spreading of viruses. Thus, you might want to block certain types of instant messages from occurring, while allowing others.

For ASA and PIX devices, IM application inspection provides detailed access control to control network usage. You can use regular expressions to help stop leakage of confidential data and the propagation of network threats. You can inspect Yahoo! Messenger or MSN Messenger traffic.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > IM (ASA 7.2+/PIX 7.2+)** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-39 Add and Edit IM Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Match Condition and Action Tab The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes, page 17-71). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the Add or Edit IM Match Criterion (for IM class maps) or Match Condition and Action (for IM policy maps) dialog boxes to do the following:

- Define the match criterion and value for an IM class map.
- Select an IM class map when creating an IM policy map.
- Define the match criterion, value, and action directly in an IM policy map.

These types of maps are used only for devices running ASA 7.2 or higher, or PIX 7.2 or higher, operating systems.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating an IM class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for IM, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating an IM policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit IM Map dialog boxes for ASA 7.2/PIX 7.2, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring IM Maps for ASA 7.2+, PIX 7.2+ Devices, page 17-70.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-40 IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing IM class map or define a new IM class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing IM class map policy object. Enter the name of the IM class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of IM traffic to match. The criteria are: • Filename—Matches the filename from IM file transfer service. • Client IP Address—Matches the source client IP address. • Client Login Name—Matches the client login name from IM service. • Peer IP Address—Matches the peer, or destination, IP address. • Peer Login Name—Matches the peer, or destination, login name from IM service. • Protocol—Matches IM protocols. • Service—Matches IM services. • File Transfer Service Version—Matches the IM file transfer service version.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.

Table 17-40 IM Class and Policy Map (ASA 7.2+/PIX 7.2+) Add or Edit Match Condition (and Action) Dialog Boxes (continued)

Element	Description
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
IP Address	The IP address you want to match.
Protocol	The IM protocol, either MSN Messenger or Yahoo! Messenger.
Services	The IM services you want to inspect. Select one or more of the listed services.

Configuring IM Maps for IOS Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX, FWSM, and IPS features/functionality, it does not support any enhancements.

Use the Add and Edit IM Map (IOS) dialog boxes to configure Instant Messaging (IM) inspection policy map objects for IOS devices. An IM map lets you change the default configuration values used for IM application inspection.

Instant Messaging causes concern due to its use of clear text when conducting business and the potential for network attacks and the spreading of viruses. Thus, you might want to block certain types of instant messages from occurring, while allowing others.

IM application inspection provides detailed access control to control network usage. It also helps stop leakage of confidential data and the propagation of network threats. The scope can be limited by identifying permitted or denied servers. Inspection of Yahoo! Messenger, MSN Messenger, and AOL instant messages are supported.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > IM (IOS)** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-41 Add and Edit IM Map (IOS) Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Service Tabs The tabs represent different IM service providers. The settings available on each tab are identical. You must configure the settings separately for each service provider. The descriptions of the following fields apply to each of the services: Yahoo!, MSN, and AOL.	
Text Chat	How you want the text chat service to be handled, for example, allowed, denied, logged, or some combination.
Other Services	How you want services other than text chat to be handled, for example, allowed, denied, logged, or some combination. IOS software recognizes all services other than text chat, such as voice-chat, video-chat, file sharing and transferring, and gaming as a single group.
Permit Servers	The servers from which to permit traffic. Accepted formats are IP addresses, IP ranges, and hostnames separated by commas.
Deny Servers	The servers from which to deny traffic. Accepted formats are IP addresses, IP ranges, and hostnames separated by commas.
Alert	Whether you want to enable or disable alerts. The default is to use the default inspection settings.
Audit	Whether you want to enable or disable an audit trail. The default is to use the default inspection settings.
Timeout	A timeout for the service. You can use the default inspection settings, or you can elect to specify a timeout. If you select Specify Timeout, enter the timeout value in seconds.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring IP Options Maps

Use the Add and Edit IP Options Map dialog boxes to define maps for the inspection of the options in an IP packet header on ASA 8.2(2)+ devices. The options field provides for control functions that are required in some situations but unnecessary for most common communications.

If you do not configure IP options inspection, the ASA device drops packets that have any options configured, with one exception. In routed mode, packets that contain the router alert option are allowed. (To disallow router alert packets, create an IP options map with router alert deselected, and configure an inspection rule to inspect IP Options using the policy map.)

Tip Because the no operation (NOP) option might be used as padding to ensure proper packet-header size and alignment, you might want to allow NOP.

For each option, you can select whether to:

- **Allow**—Allow the packet and do not change the IP header options field.
- **Clear**—Allow the packet and clear the option from the IP header options field.

If you do not select an option, the option is prohibited, and packets containing the option are dropped. Any option not listed here also results in a dropped packet; you cannot change this behavior.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > IP Options** from the Object Type selector. Right-click inside the work area, then select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-42 Add and Edit IP Options Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object.
End of Options List	End of Options List (EOOL), or IP Option 0, contains just a single zero byte and appears at the end of all options to mark the end of a list of options. This might not coincide with the end of the header according to the header length.
No operation	No Operation (NOP), or IP Option 1, is used for padding. The Options field in the IP header can contain zero, one, or more options, which makes the total length of the field variable. However, the IP header must be a multiple of 32 bits. If the number of bits of all options is not a multiple of 32 bits, the NOP option is used as to align the options on a 32-bit boundary.

Table 17-42 Add and Edit IP Options Map Dialog Boxes (continued)

Element	Description
Router alert	Router Alert (RTRALT), or IP Option 20, notifies transit routers to inspect the contents of the packet even when the packet is not destined for that router. This inspection is valuable when implementing RSVP and similar protocols require relatively complex processing from the routers along the packet's delivery path.
Basic Security *(ASA devices 9.5(1) or higher)*	IP-option Basic Security (number 130) from RFC 1108, default is to drop.
Commercial Security *(ASA devices 9.5(1) or higher)*	IP-option Commercial Security (number 134), default is to drop.
Default *(ASA devices 9.5(1) or higher)*	IP-option default configuration, default is drop.
Experimental Flow Control *(ASA devices 9.5(1) or higher)*	IP-option Experimental Flow Control (number 205), default is to drop.
Experimental Measurement *(ASA devices 9.5(1) or higher)*	IP-option Experimental Measurement (number 10), default is to drop.
Extended-Security *(ASA devices 9.5(1) or higher)*	IP-option Extended Security (number 133) from RFC 1108, default is to drop.
IMI Traffic Descriptor *(ASA devices 9.5(1) or higher)*	IP-option IMI Traffic Descriptor (number 144), default is to drop.
Quick Start *(ASA devices 9.5(1) or higher)*	IP-option Router Alert (number 25) from RFC 4782, default is to drop.
Record Route *(ASA devices 9.5(1) or higher)*	IP-option Record Route (number 7) from RFC 791, default is to drop.
Time Stamp *(ASA devices 9.5(1) or higher)*	IP-option Router Alert (number 68) from RFC 791, default is to drop.
Note Beginning with version 4.9, Security Manager supports 10 new IP Options for ASA devices running the software version 9.5(1) or higher. You can tune the inspection to allow, clear, or drop any standard or experimental options. You can also configure specific IP Options apart from the ones that are defined. For example, a value ranging between 0 and 255 can be used to configure an IP Option directly. Security Manager supports the CLI '[no] 0-255 allow | clear'. You can also set a default behavior for options not explicitly defined in an IP options inspection map. You now select which options to allow and optionally clear. For a list of IP options, with references to the relevant RFCs, see the IANA page, http://www.iana.org/assignments/ip-parameters/ip-parameters.xhtml	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 17-42 Add and Edit IP Options Map Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring IPv6 Maps

Use the Add and Edit IPv6 Map dialog boxes to define the match criteria and values for an IPv6 inspect map. You can use an IPv6 map to selectively drop IPv6 packets based on following types of extension headers found anywhere in the IPv6 packet:

- Hop-by-Hop Options
- Routing (Type 0)
- Fragment
- Destination Options
- Authentication
- Encapsulating Security Payload

Service objects corresponding to these protocols are available in the Services table in the Policy Object Manager, page 6-4.

Note With the release of Security Manager 4.4 and versions 9.0 and higher of the ASA, the separate policies for configuring IPv4 and IPv6 inspection rules were unified. However, IPv6 maps are still provided in support of earlier versions.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > IPv6** from the Object Type selector. Right-click inside the table, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-43 Add and Edit IPv6 Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.

Table 17-43 Add and Edit IPv6 Map Dialog Boxes (continued)

Element	Description
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Permit only known Extension Headers	Whether the ASA should verify the IPv6 extension header. When selected and an unknown IPv6 extension header is encountered, the ASA drops the packet and logs the action. This option is selected by default.
Enforce Extension Header Order	Whether the ASA should enforce extension header order as defined in the RFC 2460 specification. When selected and an error is detected, the ASA drops the packet and logs the action. This option is selected by default.

Match Condition and Action Tab

The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions.

These criteria entries are created and edited in the IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-78.

Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides in the Policy Object Overrides Window, page 6-20. The **Overrides** field indicates the number of devices that have overrides for this object.

IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Add or Edit Match Condition and Action dialog boxes to define an Extension Header match criterion and action for an IPv6 policy map. The contents of the Extension Headers are not processed; an action is applied based solely on the presence of a specified EH type.

The fields in these dialog boxes change based on the criterion you select.

Note You can apply multiple match definitions to one IPv6 policy map.

Navigation Path

In the Policy Object Manager, from the Match Condition and Action tab on the Add or Edit IPv6 Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring IPv6 Maps, page 17-77.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-44 IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Choose the type of IPv6 Extension Header to match: • Authentication Header (AH)—Provides integrity and data-origin authentication for IP packets. • Destination Options Header—Used for IPv6 Mobility, as well as in support of certain applications. • Encapsulating Security Payload Header (ESP)—All information following the ESP header is encrypted and not accessible to intermediate network devices. • Fragment Header—Supports traffic-source fragmented-packet communications. • Hop-by-Hop Options Header—Optional information that must be examined by every node in the packet's delivery path. • Header Count—The number of headers in the packet. When you choose this option, the following field appears; specify an upper bound for the number of headers: – Greater Than Count—Enter a value between 0 and 255. The packet is considered a match if the Header Count is greater than the specified number; it is not a match if the count is equal to, or less than the specified number. • Routing Header Type—Use this option to match one or EH types based on their header codes. When you choose this type, the following Value options appear; specify one or the other: – Routing Type—Enter one Extension Header code; for example, *51* for Authentication Header. – Routing Type Field Range—Enter a starting value and an ending value to define a range of EH codes. • Routing Header Address Count—The number of IP addresses embedded in the packet. When you choose this option, the following field appears; specify an upper bound for the number of addresses: – Greater Than Count—Enter a value between 0 and 255. The packet is considered a match if the address count is greater than the specified number; it is not a match if the count is equal to, or less than the specified number.
Type	Specifies that the map is applied only to traffic that matches the defined criteria.

Table 17-44 IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Action	Choose the action you want the device to take for traffic that matches the defined criteria: • Drop Packet—Matching packets are dropped without notification. • Drop Packet and Log—Matching packets are logged and then dropped. • Log—Matching packets are logged and processing continues.

Configuring IPsec Pass Through Maps

Use the Add and Edit IPsec Pass Through Map dialog boxes to configure settings for the IPsec Pass Through Map policy object. An IPsec Pass Through policy map lets you change the default configuration values used for IPsec Pass Through inspection.

The IPSec Pass Through inspection engine lets the security appliance pass ESP (IP protocol 50) and AH (IP protocol 51) traffic that is formed between two hosts because of successful IKE (UDP port 500) negotiation without the requirement of specific ESP or AH access lists.

The ESP or AH traffic is permitted by the inspection engine with the configured idle timeout if there is an existing control flow and it is within the connection limit defined in the MPF framework. A new control flow is created for IKE UDP port 500 traffic with the configured UDP idle timeout if there is not one, or it uses the existing flow.

To ensure that the packet arrives into the inspection engine, a hole is punched for all such traffic (ESP and AH). This inspect is attached to the control flow. The control flow is present as long as there is at least one data flow (ESP or AH) established, but the traffic always flows on the same connection. Because this IKE connection is kept open as long as data flows, a rekey would always succeed. The flows are created irrespective of whether NAT is being used. However, PAT is not supported.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > IPsec Pass Through** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-45 Add and Edit IPsec Pass Through Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.

Table 17-45 Add and Edit IPsec Pass Through Map Dialog Boxes (continued)

Element	Description
Allow ESP Maximum ESP Tunnels per Client ESP Idle Timeout	Whether to allow ESP traffic. If you select this option, you can configure the maximum number of ESP tunnels that each client can have and the amount of time that an ESP tunnel can be idle before it is closed (in hours:minutes:seconds format). The default timeout is 10 minutes (00:10:00).
Allow AH Maximum AH Tunnels per Client AH Idle Timeout	Whether to allow AH traffic. If you select this option, you can configure the maximum number of AH tunnels that each client can have and the amount of time that an AH tunnel can be idle before it is closed (in hours:minutes:seconds format). The default timeout is 10 minutes (00:10:00).
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring NetBIOS Maps

Use the Add or Edit NetBIOS Map dialog boxes to define maps for NetBIOS inspection. A NetBIOS policy map lets you change the default configuration values used for NetBIOS inspection.

The NetBIOS inspection engine translates IP addresses in the NetBIOS name service (NBNS) packets according to the security appliance NAT configuration.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > NetBIOS** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-46 Add or Edit NetBIOS Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.

Table 17-46 Add or Edit NetBIOS Map Dialog Boxes (continued)

Element	Description
Check for Protocol Violation Action	Whether to check for NETBIOS protocol violations. If you select this option, select the action you want to take when violations occur.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring ScanSafe Maps

Use the Add or Edit ScanSafe Map dialog boxes to define maps for ScanSafe inspection. A ScanSafe policy map lets you change the default configuration values used for ScanSafe inspection.

The fields on this dialog box change, depending upon whether you are creating a class map or a policy map.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > ScanSafe** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-47 ScanSafe Add Match Condition and Action Dialog Box

Element	Description
Parameters	
Transport Protocol	Allows you to select either HTTPS or HTTP. For HTTPS, the allowed range of values is 1-65535. For HTTP, the allowed range of values is 1-65535. The default value is 8080.
Default User Name	The default user name for the ScanSafe server
Default Group Name	The default group name for the ScanSafe server

Table 17-47 ScanSafe Add Match Condition and Action Dialog Box (continued)

Element	Description
Parameters	
Category	Allows you to select Cat-A through Cat-G. This is the category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.
Match Condition and Action tab only	
Class	The name of the class map
Action	Allows you to select the action you want to take when policy violations occur
+ [the "add" button]	Opens the Add Match Condition and Action dialog box. This dialog box has the following fields: • Match Type • Class Map • Action

Configuring SIP Maps

Use the Add and Edit SIP Map dialog boxes to configure values used for SIP application inspection. A SIP inspection map lets you change the default configuration values used for SIP application inspection.

SIP is a widely used protocol for Internet conferencing, telephony, presence, events notification, and instant messaging. Partially because of its text-based nature and partially because of its flexibility, SIP networks are subject to a large number of security threats.

SIP application inspection provides address translation in message header and body, dynamic opening of ports and basic sanity checks. It also supports application security and protocol conformance, which enforce the sanity of the SIP messages, as well as detect SIP-based attacks.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > SIP (ASA/PIX/FWSM)** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-48 Add and Edit SIP Map Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Enable SIP Instant Messaging Extensions	Whether to enable Instant Messaging extensions.
Permit Non-SIP Traffic on SIP Port	Whether to permit non-SIP traffic on the SIP port.
Hide Server's and Endpoint's IP Address	Whether to hide the IP addresses, which enables IP address privacy.
Check RTP Packets for Protocol Conformance Limit Payload to Audio or Video based on the Signaling Exchange	Whether to check RTP/RTCP packets flowing on the pinholes for protocol conformance. If you select this option, you can also elect to enforce the payload type to be audio/video based on the signaling exchange.
If Number of Hops to Destination is Greater Than 0	Whether to check if the value of Max-Forwards header is zero. When it is greater than zero, the action you select in the Action field is implemented. The default is to drop the packet.
If State Transition is Detected	Whether to check SIP state transitions. When a transition is detected, the action you select in the Action field is implemented. The default is to drop the packet.
If Header Fields Fail Strict Validation	Whether to take the action specified in the Action field if the SIP header fields are invalid. The default is to drop the packet.
Inspect Server's and Endpoint's Software Version	Whether to inspect the SIP endpoint software version in User-Agent and Server headers. The default is to mask the information.
If Non-SIP URI is Detected	Whether to take the action specified in the Action field if a non-SIP URI is detected in the Alert-Info and Call-Info headers. The default is to mask the information.
Match Condition and Action Tab The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see SIP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-85). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 17-48 Add and Edit SIP Map Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

SIP Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit SIP Match Criterion (for SIP class maps) or Match Condition and Action (for SIP policy maps) dialog boxes to do the following:

- Define the match criterion and value for a SIP class map.
- Select a SIP class map when creating a SIP policy map.
- Define the match criterion, value, and action directly in a SIP policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating a SIP class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for SIP, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating a SIP policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit SIP Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring SIP Maps, page 17-83.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-49 SIP Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type Class Name (Policy Map only)	Enables you to use an existing SIP class map or define a new SIP class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing SIP class map policy object. Enter the name of the SIP class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.

Table 17-49 SIP Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Specifies which criterion of SIP traffic to match. • Called Party—Matches the called party as specified in the To header. • Calling Party—Matches the calling party as specified in the From header. • Content Length—Matches the Content Length header. • Content Type—Matches the Content Type header. • IM Subscriber—Matches the SIP Instant Messenger subscriber. • Message Path—Matches the SIP Via header. • Third Party Registration—Matches the requester of a third-party registration. • URI Length—Matches a URI in the SIP headers. • Request Method—Matches the SIP request method.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Value	The regular expression you want to evaluate. You can select one of the following: • Regular Expression—The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object. • Regular Expression Group—The regular expression group object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression group object.
URI Type	The type of URI to match, either SIP or TEL.
Greater Than Length	The length in bytes of the evaluated field. The criterion matches if the length is greater than the specified number, and does not match if the field is less than the specified number.

Table 17-49 SIP Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Content Type	The content type to evaluate as specified in the content-type header field. You can select one of the following: • SDP—Matches an SDP SIP content header type. • Regular Expression, Regular Expression Group—The regular expression or regular expression group to evaluate. See the explanation for the Value field for an explanation of these options.
Resource Method	The request method you want to inspect: • ack—Confirms that the client has received a final response to an INVITE request. • bye—Terminates a call and can be sent by either the caller or the called party. • cancel—Cancels any pending searches but does not terminate a call that has already been accepted. • info—Communicates mid-session signaling information along the signaling path for the call. • invite—Indicates a user or service is being invited to participate in a call session. • message—Sends instant messages where each message is independent of any other message. • notify—Notifies a SIP node that an event which has been requested by an earlier SUBSCRIBE method has occurred. • options—Queries the capabilities of servers. • prack—Provisional response acknowledgment. • refer—Requests that the recipient REFER to a resource provided in the request. • register—Registers the address listed in the To header field with a SIP server. • subscribe—Requests notification of an event or set of events at a later time. • unknown—Uses a nonstandard extension that could have unknown security impacts on the network. • update—Permits a client to update parameters of a session but has no impact on the state of a dialog.

Configuring Skinny Maps

Use the Add or Edit Skinny Map dialog boxes to define Skinny maps for Skinny inspection. A Skinny policy map lets you change the default configuration values used for Skinny inspection.

Skinny (SCCP) is a simplified protocol used in VoIP networks. Cisco IP Phones using SCCP can coexist in an H.323 environment. When used with Cisco CallManager, the SCCP client can interoperate with H.323 compliant terminals. Application layer functions in the security appliance recognize SCCP version 3.3. There are 5 versions of the SCCP protocol: 2.4, 3.0.4, 3.1.1, 3.2, and 3.3.2.

The security appliance supports all versions through 3.3.2. The security appliance supports PAT and NAT for SCCP. PAT is necessary if you have more IP phones than global IP addresses for the IP phones to use. By supporting NAT and PAT of SCCP Signaling packets, Skinny application inspection ensures that all SCCP signaling and media packets can traverse the security appliance.

Normal traffic between Cisco CallManager and Cisco IP Phones uses SCCP and is handled by SCCP inspection without any special configuration. The security appliance also supports DHCP options 150 and 66, which it accomplishes by sending the location of a TFTP server to Cisco IP Phones and other DHCP clients. Cisco IP Phones might also include DHCP option 3 in their requests, which sets the default route.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > Skinny** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-50 Add and Edit Skinny Map Dialog Boxes

Element	Description
Name	The name of the Skinny map. A maximum of 40 characters is allowed.
Description	A description of the Skinny map, up to 200 characters.
Parameters Tab	
Enforce Endpoint Registration	Whether to enforce registration before calls can be placed.
Maximum SCCP Station Message ID 0x	The maximum SCCP station message ID allowed, in hexadecimal.
Check RTP Packets for Protocol Conformance Enforce Payload Type to be Audio or Video based on Signaling Exchange	Whether to check RTP packets flowing through the pinholes for protocol conformance. If you select this option, you can also select whether to enforce the payload type.
Minimum SCCP Prefix Length	The minimum SCCP length allowed.
Maximum SCCP Prefix Length	The maximum SCCP length allowed.
Media Timeout	The timeout value for media connections.
Signaling Timeout	The timeout value for signaling connections.

Table 17-50 Add and Edit Skinny Map Dialog Boxes (continued)

Element	Description
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see Skinny Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-89). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Skinny Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Add or Edit Match Condition and Action dialog boxes to define the match criterion, value, and action for a Skinny policy map.

Navigation Path

In the Policy Object Manager, from the Match Condition and Action tab on the Add or Edit Skinny Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring SIP Maps, page 17-83.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-51 Skinny Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Specifies which criterion of Skinny traffic to match.
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on 0xFFFF, then any traffic that has the message ID 0xFFFF is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.

Table 17-51 Skinny Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
ID Type	The hexadecimal value for the message ID to inspect: • Value—Matches a single hexadecimal value. • Range—Matches a range of values.
Action	The action you want the device to take for traffic that matches the defined criteria.

Configuring SNMP Maps

Use the Add and Edit SNMP Map dialog boxes to define maps for SNMP inspection. An SNMP policy map lets you change the default configuration values used for SNMP application inspection.

SNMP application inspection lets you restrict SNMP traffic to a specific version of SNMP. Earlier versions of SNMP are less secure; therefore, denying certain SNMP versions may be required by your security policy. The security appliance can deny SNMP versions 1, 2, 2c, or 3. You control the versions permitted by creating an SNMP map. You then apply the SNMP map when you enable SNMP inspection.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > SNMP** from the Object Type selector. Right-click inside the work area, then select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-52 Add and Edit SNMP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Disallowed SNMP Versions	The versions of SNMP you want to prohibit. • SNMP Version 1 • SNMP Version 2c (Community Based) • SNMP Version 2 (Party Based) • SNMP Version 3
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 17-52 Add and Edit SNMP Map Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring SCTP Maps

SCTP is a transport-layer protocol operating on top of IP in the protocol stack, similar to TCP and UDP. SCTP creates a logical communication channel, called an association, between two end nodes over multiple source or destination IP addresses. An association defines a set of IP addresses on each node (source and destination) and a port on each node. Any IP address can be used as either a source or a destination IP address of data packets in the association. Messages can be transmitted between a pair of IP addresses, which is defined as a stream.

If you have SCTP traffic going through the ASA, you can configure Cisco Security Manager to control access based on SCTP ports, and implement application layer inspection to enable connections and to optionally filter on payload protocol ID (PPID) to selectively drop, log, or rate limit applications.

You can refine your access rules by adding an SCTP inspect map and filtering on SCTP applications.You can selectively drop, log, or rate limit SCTP traffic classes based on the payload protocol identifier (PPID).

When you filter on PPID, keep the following in mind:

- PPIDs are in data chunks, and a given packet can have multiple data chunks. If a packet includes data chunks with different PPIDs, the packet will not be filtered, and the assigned action will not be applied to the packet.
- If you use PPID filtering to drop or rate-limit packets, be aware that the transmitter will resend any dropped packets. Although a packet for a rate-limited PPID might make it through on the next attempt, a packet for a dropped PPID will again be dropped. You might want to evaluate the eventual consequence of these repeated drops on your network.

Use the Add and Edit SCTP Map dialog boxes to define the match criteria and values for an SCTP inspect map. You can use an SCTP map to inspect packets based on the Payload PID criteria. You can perform the following actions on the packets, based on the PPID match criteria:

- No Action
- Drop Packet
- Log Packet
- Rate Limit

Service objects corresponding to the SCTP protocol are available in the Services table in the Policy Object Manager, page 6-4.

Note SCTP inspect maps are supported from Security Manager 4.10 and ASA versions 9.5.2 and higher.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > SCTP** from the Object Type selector. Right-click inside the table, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-53 Add and Edit SCTP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. These criteria entries are created and edited in the SCTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-92.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides in the Policy Object Overrides Window, page 6-20. The **Overrides** field indicates the number of devices that have overrides for this object.

SCTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Add or Edit Match Condition and Action dialog boxes to define a Payload PID match criterion and action for a SCTP policy map. Repeat the process until you identify all PPIDs you want to selectively handle.

Navigation Path

In the Policy Object Manager, from the Match Condition and Action tab on the Add or Edit IPv6 Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring IPv6 Maps, page 17-77.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-54 IPv6 Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Select the Payload PID (PPID) criterion.
Type	Specifies that the map is applied only to traffic that matches or does not match the defined criteria.
You can find the current list of SCTP PPIDs at http://www.iana.org/assignments/sctp-parameters/sctp-parameters.xhtml#sctp-parameters-25.	
Min. Payload PID	Enter a PPID number. There are certain PPIDs associated with a name, which Cisco Security Manager accepts, and processes internally. Enter the PPID number in the text box, and click OK. The corresponding name will be displayed in the match action table if it matches the default names.
Max. Payload PID	(Optional) Enter a second, higher PPID to specify a range of PPIDs.
Action	Choose the action based on the PPID in SCTP data chunks: • Drop Packet—Drop and log all packets that match. • Log—Send a system log message. • Rate Limit— Limit the rate of messages. The rate is in packets per second.

Configuring Diameter Maps

Diameter is an Authentication, Authorization, and Accounting (AAA) protocol used in next-generation mobile and fixed telecom networks such as EPS (Evolved Packet System) for LTE (Long Term Evolution) and IMS (IP Multimedia Subsystem). It replaces RADIUS and TACACS in these networks.

Diameter uses TCP and SCTP as the transport layer, and secures communications using TCP/TLS and SCTP/DTLS. It can optionally provide data object encryption as well. For detailed information on Diameter, see RFC 6733.

Diameter applications perform service management tasks such as deciding user access, service authorization, quality of service, and rate of charging. Although Diameter applications can appear on many different control-plane interfaces in the LTE architecture, the ASA inspects Diameter command codes and attribute-value pairs (AVP) for the following interfaces only:

- S6a: Mobility Management Entity (MME) - Home Subscription Service (HSS).
- S9: PDN Gateway (PDG) - 3GPP AAA Proxy/Server.
- Rx: Policy Charging Rules Function (PCRF) - Call Session Control Function (CSCF).

Diameter inspection opens pinholes for Diameter endpoints to allow communication. The inspection supports 3GPP version 12 and is RFC 6733 compliant.

You can use the Add and Edit Diameter Map dialog boxes to filter traffic based on application ID, command codes, and AVP, to apply special actions such as dropping packets or connections, or logging them. You can create custom AVP for newly-registered Diameter applications. Filtering lets you fine-tune the traffic you allow on your network. For more information see Create and Add Custom AVPs, page 17-97

Note Diameter messages for applications that run on other interfaces will be allowed and passed through by default. However, you can configure a Diameter inspection policy map to drop these applications by application ID, although you can specify actions based on the command codes or AVP for these unsupported applications.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > Diameter** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28
- Create and Add Custom AVPs, page 17-97

Field Reference

Table 17-55 Add and Edit Diameter Map Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Unsupported application-id action log	To log unsupported Diameter application identifier (Diameter application name) in the map. Application ID is a number between 0-4294967295, in the map. These applications are registered with the IANA. Following are the core supported applications, but you can filter on other applications. **3gpp-rx-ts29214** (16777236) **3gpp-s6a** (16777251) **3gpp-s9** (16777267) **common-message** (0) - This is the base Diameter protocol
Unsupported command code action log	To log unsupported Diameter command codes in the map, where *code* is the Diameter command code name or number (0-4294967295).
Unsupported avp action log	To log unsupported attribute- value pair parameter
Strict Parameters	
Enable Session Validation	To validate session-ID AVP related messages
Enable State Validation	To enable validation of state machine

Table 17-55 Add and Edit Diameter Map Dialog Box (continued)

Element	Description
Match Condition and Action Tab	
	The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-95). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes

Use the Add or Edit Diameter Match Criterion (for Diameter class maps) or Match Condition and Action (for Diameter policy maps) dialog boxes to do the following:

- Define the match criterion and value for a Diameter class map.
- Select a Diameter class map when creating a Diameter policy map.
- Define the match criterion, value, and action directly in a Diameter policy map.

The fields on this dialog box change based on the criterion you select and whether you are creating a class map or policy map.

Navigation Path

When creating a Diameter class map, in the Policy Object Manager, from the Add or Edit Class Maps dialog boxes for Diameter, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Class Maps for Inspection Policies, page 17-28.

When creating a Diameter policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit Diameter Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring Diameter Maps, page 17-93.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-56 Diameter Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Match Type (Only Policy Map)	Enables you to use an existing Diameter class map or define a new Diameter class map. • Use Specified Values—You want to define the class map on this dialog box. • Use Values in Class Map—You want to select an existing Diameter class map policy object. Enter the name of the Diameter class map in the **Class Name** field. Click **Select** to select the map from a list or to create a new class map object.
Criterion	Specifies which criterion of Diameter traffic to match. • Application ID—Matches the application identifier, where the application identifier is the Diameter application name or number (0-4294967295) in the **Begin Value** field. If there is a range of consecutively-numbered applications that you want to match, you can include a second ID in the **End Value** field. You can define the range by application name or number, and it applies to all the numbers between the **Begin Value** and the **End Value**. • Command Code—Matches the command code, where *code* is the Diameter command code name or number (0-4294967295) in the **Begin Value** field. If there is a range of consecutively-numbered command codes that you want to match, you can include a second code in the **End Value** field. You can define the range by command code name or number, and it applies to all the numbers between the **Begin Value** and the **End Value**. • AVP—Matches the Attribute Value Pair. – To match AVP based on attribute only, specify the name or number (1-4294967295) of an attribute-value pair. For the first code, you can specify the name of a custom AVP or one that is registered in RFCs or 3GPP technical specifications and is directly supported in the software in the **Begin Value** field. If you want to match a range of AVP, specify the second code by number only in the **End Value** field. If you want to match an AVP by its value, you cannot specify a second code. Specify the ID number of the vendor to also match, from 0-4294967295 in the **Vendor ID** field. For example, the 3GPP vendor ID is 10415, the IETF is 0. – To match AVP based on the value of the attribute, additionally specify the value of the attribute in the **AVP Data Type** field. **Note** You can create and add custom AVPs to new diameter applications. For more information see, Create and Add Custom AVPs, page 17-97

Table 17-56 Diameter Class and Policy Maps Add and Edit Match Condition and Action Dialog Boxes (continued)

Element	Description
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
AVP DataType	You can configure this only if the data type of the AVP is supported. For example, you can specify an IP address for AVP that have the address data type. Following are the specific syntax of the value option for the supported data types • Address—Specify the IPv4 or IPv6 address to match. For example, 10.100.10.10 or 2001:DB8::0DB8:800:200C:417A • Diameter Identity, Diameter URI, Octet String, UTF8tString—Use regular expression or regular expression class objects to match these data types. • Enumerated—Specify a range of numbers in the Begin Range and End Range fields. The range is 0 - 4294967295. • Float32: decimal point representation with 8 digit precision • Float64: decimal point representation with 16 digit precision • Integer32: -2147483647 to 2147483647 • Integer64: -9223372036854775807 to 9223372036854775807 • Unsigned32: 0 to 4294967295 • Unsigned64: 0 to 18446744073709551615 • Time— Specify the start and end dates and time. Both are required. Time is in 24-hour format. **Note** You can create and add custom AVPs to new Diameter applications.
Action (Policy Map only)	The action you want the device to take for traffic that matches the defined criteria.

Create and Add Custom AVPs

Use the Add AVP dialog boxes to create and add custom AVPs. These can be registered with the IETF and added to new Diameter applications.

Note Cisco Security Manager does not allow you to edit a custom AVP object, once created.However the Device Override option allows you to edit the custom AVP for a particular device. If you want to change any parameter in the custom AVP object, you have to remove the custom AVP reference from the diameter building block (if it is referred), deploy to the device (if it is present in the device) and re-create the object with the required values and refer it back in the diameter building block and deploy it again.

Navigation Path

When creating a custom AVP, in the Policy Object Manager, from the Add Match Criterion dialog box for Diameter, select **AVP in the Criterion**, then select **Begin Value and right click in the AVP Maps Selector dialog box to Add AVP.**

Field Reference

Table 17-57 Add AVP Dialog Boxes

Element	Description
Name	The name of the custom AVP. A maximum of 32 characters is allowed. **Note** At least one character of the name must be an alphabet.
Description	A description of the AVP. A maximum of 80 characters is allowed.
AVP Code	Set a value for the AVP Code (256- 4294967295), that belongs to the specific vendor code address space.
DataType	You can configure this only if the data type of the AVP is supported. For example, you can specify an IP address for AVP that have the address data type. Following are the specific syntax of the value option for the supported data types • Address—Specify the IPv4 or IPv6 address to match. For example, 10.100.10.10 or 2001:DB8::0DB8:800:200C:417A • Diameter Identity, Diameter URI, Octet String, UTF8tString—Use regular expression or regular expression class objects to match these data types. • Enumerated—Specify a range of numbers in the Begin Range and End Range fields. The range is 0 - 4294967295. • Float32: decimal point representation with 8 digit precision • Float64: decimal point representation with 16 digit precision • Integer32: -2147483647 to 2147483647 • Integer64: -9223372036854775807 to 9223372036854775807 • Unsigned32: 0 to 4294967295 • Unsigned64: 0 to 18446744073709551615 • Time— Specify the start and end dates and time. Both are required. Time is in 24-hour format.
Vendor ID	Specify the ID number of the vendor, from 0-4294967295 in the **Vendor ID** field. For example, the 3GPP vendor ID is 10415, the IETF is 0.

Table 17-57 Add AVP Dialog Boxes (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Create and Add TLS Proxy Objects

If a Diameter application uses encrypted data over TCP, inspection cannot see inside the packets to implement your message filtering rules. Thus, if you create filtering rules, and you want them to also apply to encrypted TCP traffic, you must configure a TLS proxy. You also need a proxy if you want strict protocol enforcement on encrypted traffic. This configuration does not apply to SCTP/DTLS traffic.

The TLS proxy acts as a man-in-the-middle. It decrypts traffic, inspects it, then encrypts it again and sends it to the intended destination. Thus, both sides of the connection, the Diameter server and Diameter client, must trust the ASA, and all parties must have the required certificates. You must have a good understanding of digital certificates to implement TLS proxy.

Note The TLS proxy feature is supported in multi-context devices for version ASA 9.7.1 and higher.

You have the following options for configuring TLS proxy for Diameter inspection:

- Full TLS proxy—Encrypt traffic between the ASA and Diameter clients and the ASA and Diameter server. You have the following options for establishing the trust relationship with the server:
  - Use a static proxy client trustpoint. The ASA presents the same certificate for every Diameter client when communicating with the Diameter server. Because all clients look the same, the Diameter server cannot provide differential services per client. On the other hand, this option is faster than the LDC method.
  - Use local dynamic certificates (LDC). With this option, the ASA presents unique certificates per Diameter client when communicating with the Diameter server. This method gives the Diameter server better visibility into client traffic, which makes it possible to provide differential services based on client characteristics.
- TLS offload—Encrypt traffic between the ASA and Diameter client, but use a clear-text connection between the ASA and Diameter server. This option is viable if the Diameter server is in the same data center as the ASA, where you are certain that the traffic between the devices will not leave the protected area. Using TLS offload can improve performance, because it reduces the amount of encryption processing required. It should be the fastest of the options. The Diameter server can apply differential services based on client IP address only.

Navigation Path

Select **Manage > Policy Objects**, then select **TLS Proxy** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object.**

Field Reference

Table 17-58 Add TLS Proxy Dialog Boxes

Element	Description
Name	The name of the TLS Proxy object. A maximum of 63 characters is allowed. **Note** At least one character of the name must be an alphabet.
Description	A description of the TLS Proxy object.
Server Configuration	
Server Proxy Certificate	Click **Select** to import the CA certificate that is used to sign the Diameter client's certificate into an ASA trustpoint. This step specifies the proxy trustpoint certificate to be presented during TLS handshake. The trustpoint could be self-signed or issued by a third party. This allows the ASA to trust the Diameter clients.
Enable client authentication during TLS proxy handshake	Select to require the ASA to present a certificate and authenticate the TLS client during TLS handshake.
Encryption (Optional)	Beginning with 4.14, Cisco Security Manager allows you to configure cipher suites, when TLS Proxy is used as server. This field defines the cipher suites to be announced/matched during the TLS handshake. Select the Hashing algorithms, which are needed for encryption of data, from the Available Members List and add them to Selected Members list.
Client Configuration	
Configure the proxy client to use clear text to communicate with the remote TCP server	Select proxy client to use clear text, if encryption is not needed.
Specify the proxy certificate for the TLS client. The client proxy certificate could either be self-signed, enrolled with a CA or issued by a third party.	Select to specify Client Proxy Certificate. Alternately, click **Select** to import the CA certificate for the TLS client.
Specify the internal Certificate Authority to sign the local dynamic certificates for phones. This local CA can be self-signed certificate with proxy-ldc-issuer enabled or you may use embedded Local CA Server to issue LDC to phones.	Select to specify Local Dynamic Certificate Issuer. Alternately, click **Select** to import the CA certificate, which could serve as Local Dynamic certificate (LDC) issuer

Table 17-58 Add TLS Proxy Dialog Boxes (continued)

Element	Description
Local Dynamic Certificate Key Pair	
Key Pair Name	Specifies the RSA key pair to be used by the client or server's dynamic certificates. The key pair must have been generated with the "crypto key generate" command. The keypair must exist on the device before deployment.
Encryption (Optional)	Defines the cipher suites to be announced/matched during the TLS handshake. For client proxy (the proxy acts as a TLS client to the server), the user-defined cipher suites replace the original ones from the Hello message for asymmetric encryption method between the two TLS legs. Select the Hashing algorithms, which are needed for encryption of data, from the Available Members List and add them to Selected Members list. **Note** From ASA version 9.7.1, the Cisco Security Manager supports TLS1.2 new cipher suites— aes256-sha384 and aes128-sha256.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Editing TLS Proxy Object

Cisco Security Manager does not allow you to edit a TLS proxy object, once created. However, the Device Override option allows you to edit the TLS proxy object for a particular device.

If you want to change any parameter in the TLS proxy object, you have to remove the TLS proxy reference from the diameter building block (if it is referred), deploy to the device (if it is present in the device) and re-create the object with a new name, with the required values and refer it back in the diameter building block and deploy it again.

To edit a TLS proxy in class-map, execute the following deployment procedure:

1. Remove the relevant class map with existing TLS proxy server from the device by navigating to **Platform > Service Policy > Rules**.
2. Deploy the relevant class map with new TLS proxy server to the device by navigating to **Platform > Service Policy > Rules**.

Configuring LISP Maps

The Locator ID Separation Protocol (LISP) is a network architecture and protocol. LISP replaces a single IP address with two numbering spaces—Routing Locators (RLOCs), which are topologically assigned to network attachment points and used for routing and forwarding of packets through the network; and Endpoint Identifiers (EIDs), which are assigned independently from the network topology and used for numbering devices, and are aggregated along administrative boundaries.

LISP defines functions for mapping between the two numbering spaces and encapsulating traffic originated by devices using non-routable EIDs for transport across a network infrastructure that routes and forwards using RLOCs. LISP provides a set of functions for devices to exchange information that is used to map non-routable EIDs to routable RLOCs.

When considering the deployment of ACLs with LISP, the following aspects are important.

- LISP encapsulation utilizes a UDP header just prior to the LISP header for all packets to distinguish between two distinct packet groups: LISP control plane packets, which utilize a UDP destination port of 4342, and LISP data plane packets, which utilize a UDP destination port of 4341. ACLs may need to consider this distinction between these two groups of packets.
- LISP is an encapsulation protocol and, because ACLs only filter based on Layer 3 and Layer 4 header information, ACLs may need to be applied at a specific point or at several different points within the packet forwarding and LISP encapsulation process in order to implement a site security policy. The application point and direction of the ACL will dictate whether EID namespace or RLOC namespace is used within the ACL itself. Packets can be filtered using EID namespace just prior to LISP encapsulation or just after LISP decapsulation; packets can be filtered using RLOC namespace just after LISP encapsulation or just prior to LISP decapsulation.

You can use the Add and Edit LISP Map dialog boxes to filter traffic based on EID access-list and validation key. Filtering lets you fine-tune the traffic you allow on your network.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > LISP** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-59 Add and Edit LISP Map Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Allowed Eid access-list	Enables you to select a unified access list building block.
Validation key	Specify an unencrypted clear text password.

Table 17-59 Add and Edit LISP Map Dialog Box (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring M3UA Maps

MTP3 User Adaptation (M3UA) is a client/server protocol that provides a gateway to the SS7 network for IP-based applications that interface with the SS7 Message Transfer Part 3 (MTP3) layer. M3UA makes it possible to run the SS7 User Parts (such as ISUP) over an IP network. M3UA is defined in RFC 4666.

M3UA uses SCTP as the transport layer. SCTP port 2905 is the expected port, although you can configure the signaling gateways to use a different port.

The MTP3 layer provides networking functions such as routing and node addressing, but uses point codes to identify nodes. The M3UA layer exchanges Originating Point Codes (OPC) and Destination Point Codes (DPC). This is similar to how IP uses IP addresses to identify nodes.

M3UA inspection provides limited protocol conformance. You can optionally apply access policy based on point codes or Service Indicators (SI). You can also apply rate limiting based on message class and type.

M3UA Protocol Conformance

M3UA inspection provides the following limited protocol enforcement. Inspection drops and logs packets that do not meet requirements.

- Common message header. Inspection validates all fields in the common header.
  - Version 1 only.
  - Message length must be correct.
  - Message type class with a reserved value is not allowed.
  - Invalid message ID within the message class is not allowed.
- Payload data message.
  - Only one parameter of a given type is allowed.
  - Data messages on SCTP stream 0 are not allowed.

M3UA Inspection Limitations

M3UA inspection has the following limitations.

- NAT is not supported for IP addresses embedded in M3UA data.
- Segmented M3UA messages will not be inspected and are likely to be dropped.
- SCTP does not support multi-homing or multi-streaming. If you need to support multi-homed flows you need to create access lists to allow them.
- Stateful failover is not supported for call flows and messages. Any failure occurring during a call flow might cause packets to be dropped and calls to be disconnected.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Inspect > M3UA** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Configuring Class Maps for Inspection Policies, page 17-28

Field Reference

Table 17-60 Add and Edit M3UA Map Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
SS7 Variant	Select the SS7 variant that will be used in your network for M3UA inspection.This option determines the valid format for point codes. After you configure the option and deploy an M3UA policy, you cannot change it unless you first remove the policy. The default SS7 variant is ITU.
Enable M3UA Application Server Process (ASP) State Validation	Select to perform application server process (ASP) state validation. The system maintains the ASP states of M3UA sessions and allows or drops ASP messages based on the validation result. If you do not enable strict ASP state validation, all ASP messages are forwarded uninspected.
Enforce Timeout	
Endpoint	Enter the idle timeout to remove statistics for an M3UA endpoint, in the hh:mm:ss format. To have no timeout, specify 0. The default is 30 minutes (00:30:00).
Session	Enter the idle timeout to remove an M3UA session if you enable strict ASP state validation, in hh:mm:ss format. To have no timeout, specify 0. The default value is 30 minutes (00:30:00). When this timeout is disabled, the system cannot remove stale sessions.

Table 17-60 Add and Edit M3UA Map Dialog Box (continued)

Element	Description
M3UA Message Tag Validation	
Specify, whether to check and validate the content of certain fields for the specified message type. Messages that fail validation are dropped. Validation differs by message type. Select the messages you want to validate.	
Destination User Part Unavailable (DUPU)	The User/Cause field must be present, and it must contain only valid cause and user codes.
Error	All mandatory fields must be present and must contain only allowed values. Each error message must contain the required fields for that error code.
Notify	The status type and status information fields must contain allowed values only.
Match Condition and Action Tab	
The Match All table lists the criteria included in the policy map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion, the criterion and value that is inspected, and the action to be taken for traffic that satisfies the conditions. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see M3UA Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-105). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

M3UA Policy Maps Add or Edit Match Condition and Action Dialog Boxes

Use the Match Condition and Action dialog boxes to define the match criterion, value, and action directly in a M3UA policy map.

The fields on this dialog box change based on the criterion you select while creating a policy map.

Navigation Path

When creating a M3UA policy map, in the Policy Object Manager, from the Match Condition and Action tab on the Add and Edit M3UA Map dialog boxes, right-click inside the table, then select **Add Row** or right-click a row, then select **Edit Row**. See Configuring M3UA Maps, page 17-103.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22

Field Reference

Table 17-61 M3UA Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Criterion	Specifies which criterion of SCTP traffic to match - Message, DPC, OPC, or Service Indicator.
Message criterion	Matches the M3UA message class and type. The possible values for message class ID and its corresponding message ID are detailed here. Refer to the M3UA RFCs and documentation for detailed information about these messages. • class ID 0 (Management Messages)- message ID 0-1 • class ID 1(Transfer Messages)- message ID 1 • class ID 2(SS7 Signaling Network Management Messages)- message ID 1-6 • class ID 3(ASP State Maintenance Messages)-message ID 1-6 • class ID 4(ASP Traffic Maintenance Messages)- message ID 1-4 • class ID 9(Routing Key Management Messages)- message ID 1-4
DPC criterion	Matches the destination point code in the data message. Point code is in the zone-region-sp format, where the possible values for each element depend on the SS7 variant.
OPC criterion	Matches the originating point code in the data message, that is, the traffic source. Point code is in zone-region-sp format, where the possible values for each element depend on the SS7 variant: • ITU—Point codes are 14 bit in 3-8-3 format. The value ranges are [0-7]-[0-255]-[0-7]. • ANSI—Point codes are 24 bit in 8-8-8 format. The value ranges are [0-255]-[0-255]-[0-255]. • Japan—Point codes are 16 bit in 5-4-7 format. The value ranges are [0-31]-[0-15]-[0-127]. • China—Point codes are 24 bit in 8-8-8 format. The value ranges are [0-255]-[0-255]-[0-255].

Table 17-61 M3UA Policy Maps Add and Edit Match Condition and Action Dialog Boxes

Element	Description
Service Indicator criterion	Matches the service indicator number, 0-15. The available service indicators are listed in the variables section. Consult M3UA RFCs and documentation for detailed information about these service indicators • 0—Signaling Network Management Messages • 1—Signaling Network Testing and Maintenance Messages • 2—Signaling Network Testing and Maintenance Special Messages • 3—SCCP • 4—Telephone User Part • 5—ISDN User Part • 6—Data User Part (call and circuit-related messages) • 7—Data User Part (facility registration and cancellation messages) • 8—Reserved for MTP Testing User Part • 9—Broadband ISDN User Part • 10—Satellite ISDN User Part • 11—Reserved • 12—AAL type 2 Signaling • 13—Bearer Independent Call Control • 14—Gateway Control Protocol • 15—Reserved
Type	Specifies whether the map includes traffic that matches or does not match the criterion. For example, if Doesn't Match is selected on the string "example.com," then any traffic that contains "example.com" is excluded from the map. • Matches—Matches the criterion. • Doesn't Match—Does not match the criterion.
Action	The action you want the device to take for traffic that matches the defined criteria. • Drop Packet—By default, all invalid packets or packets that failed during parsing are dropped. • Drop Packet and Log— Same as drop packet and additionally send a system log message. • Rate Limit— Limit the rate of messages. This option is available when the message criterion is selected.

Configuring Regular Expression Groups

Use the Add and Edit Regular Expression Groups dialog boxes to define regular expression groups, which contain multiple regular expressions. Groups make it possible for you to create modular regular expressions and group them in multiple ways for various uses. The objects can be used in some inspection class maps and inspection policy maps.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Regular Expressions Groups** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Creating Policy Objects, page 6-9

Field Reference

Table 17-62 Add and Edit Regular Expression Class Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Regular Expressions	The Regular Expression policy objects that include the expressions you want to include in the group. Enter the name of the objects or click **Select** to select them from a list or to create a new object.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add/Edit Regular Expressions

Use the Add and Edit Regular Expression dialog boxes to define regular expressions for use in class and policy inspection maps or in regular expression group policy objects. Regular expressions are also used in remote access SSL VPN client settings.

A regular expression matches text strings either literally as an exact string or by using metacharacters so you can match multiple variants of a text string. You can use regular expressions in various type of class and policy inspection maps to match various target items, for example, the content of certain application traffic such as the body text inside an HTTP packet.

Navigation Path

- Select **Manage > Policy Objects**, then select **Maps > Regular Expressions** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.
- From the Client Settings tab of the SSL VPN Other Settings policy for ASA devices, click the **Add Row** button for the AnyConnect Client Image table, or select an image and click the **Edit Row** button. For detailed information on opening the tab, see Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64. On the Add AnyConnect Client Image dialog box, click **Select** to open the Regular Expressions Selector dialog box. To add a new regular expression, click the **Add (+)** button on the Regular Expressions Selector dialog box.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Protocols and Maps for Inspection, page 17-22
- Creating Policy Objects, page 6-9

Field Reference

Table 17-63 Add and Edit Regular Expression Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Value	The regular expression, up to 100 characters in length. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions, page 17-109.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Metacharacters Used to Build Regular Expressions

The following table explains the metacharacters you can use to build regular expressions in the Add and Edit Regular Expression dialog boxes (see Add/Edit Regular Expressions, page 17-108).

Keep the following tips in mind when creating regular expressions:

- If you enter any metacharacters in your text string that you want to be used literally, add the backslash (\) escape character before them. For example, “example\.com”.
- If you want to match upper and lower case characters, enter text in both upper- and lowercase. For example, “cats” is entered as “[cC][aA][tT][sS]”.

Table 17-64 Metacharacters Used to Build Regular Expressions

Character	Description	Notes
.	Dot	Matches any single character. For example, d.g matches dog, dag, dtg, and any word that contains those characters, such as doggonnit.
(exp)	Subexpression	A subexpression segregates characters from surrounding characters, so that you can use other metacharacters on the subexpression. For example, d(ola)g matches dog and dag, but dolag matches do and ag. A subexpression can also be used with repeat quantifiers to differentiate the characters meant for repetition. For example, ab(xy){3}z matches abxyxyxyz.
|	Alternation	Matches either expression it separates. For example, dog|cat matches dog or cat.
?	Question mark	A quantifier that indicates that there are 0 or 1 of the previous expression. For example, lo?se matches lse or lose.
*	Asterisk	A quantifier that indicates that there are 0, 1 or any number of the previous expression. For example, lo*se matches lse, lose, loose, etc.
+	Plus	A quantifier that indicates that there is at least 1 of the previous expression. For example, lo+se matches lose and loose, but not lse.
{x}	Repeat Quantifier	Repeat exactly x times. For example, ab(xy){3}z matches abxyxyxyz.
	Minimum repeat quantifier	Repeat at least x times. For example, ab(xy){2,}z matches abxyxyz, abxyxyxyz, etc.
[abc]	Character class	Matches any character in the brackets. For example, [abc] matches a, b, or c.
[^abc]	Negated character class	Matches a single character that is not contained within the brackets. For example, [^abc] matches any character other than a, b, or c. [^A-Z] matches any single character that is not an uppercase letter.
[a-c]	Character range class	Matches any character in the range. [a-z] matches any lowercase letter. You can mix characters and ranges: [abcq-z] matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does [a-cq-z]. The dash (-) character is literal only if it is the last or the first character within the brackets: [abc-] or [-abc].
“”	Quotation marks	Preserves trailing or leading spaces in the string. For example, “ test” preserves the leading space when it looks for a match.
^	Caret	Specifies the beginning of a line.
\	Escape character	When used with a metacharacter, matches a literal character. For example, \[matches the left square bracket.

Table 17-64 Metacharacters Used to Build Regular Expressions (continued)

Character	Description	Notes
char	Character	When character is not a metacharacter, matches the literal character.
\r	Carriage return	Matches a carriage return 0x0d.
\n	Newline	Matches a new line 0x0a.
\t	Tab	Matches a tab 0x09.
\f	Formfeed	Matches a form feed 0x0c.
\xNN	Escaped hexadecimal number	Matches an ASCII character using hexadecimal (exactly two digits).
\NNN	Escaped octal number	Matches an ASCII character as octal (exactly three digits). For example, the character 040 represents a space.

Configuring Settings for Inspection Rules for IOS Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX, FWSM, and IPS features/functionality, it does not support any enhancements.

If you configure inspection rules, you can also configure inspection settings to change the default settings for some global inspection parameters for IOS devices. Most of the inspection settings relate to preventing or mitigating Denial of Service (DoS) attacks. The default settings for most of these options are appropriate for most networks, so configure this policy only if you need to adjust one or more settings. If you do not change a setting, it is not configured on the device (the default remains configured).

To open the Inspection settings page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Inspection** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Inspection** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Settings > Inspection**.

The following table explains the available inspection settings.

Table 17-65 Inspection Page

Element	Description
Global Timeout Values	
TCP Establish Timeout (seconds)	How long to wait for a TCP session to reach the established state before dropping the session, in seconds, from 1-2147483. The default is 30.
FIN Wait Time (seconds)	How long to maintain TCP session state information after the firewall detects a FIN-exchange, in seconds, from 1-2147483. The FIN-exchange occurs when the TCP session is ready to close. The default is 5.

Table 17-65 Inspection Page (continued)

Element	Description
TCP Idle Time (seconds)	How long to maintain a TCP session while there is no activity in the session, in seconds, from 1-2147483. The default is 3600 (one hour).
UDP Idle Time (seconds)	How long to maintain a UDP session while there is no activity in the session, in seconds, from 1-2147483. The default is 30. When the software detects a valid UDP packet, the software establishes state information for a new UDP session. Because UDP is a connectionless service, there are no actual sessions, so the software approximates sessions by examining the information in the packet and determining if the packet is similar to other UDP packets (for example, it has similar source or destination addresses) and if the packet was detected soon after another similar UDP packet. If the software detects no UDP packets for the UDP session for the period of time defined by the UDP idle timeout, the software will not continue to manage state information for the session.
DNS Timeout (seconds)	The length of time for which a DNS lookup session is managed while there is no activity, in seconds, from 1-2147483. The default is 5.
SYN Flooding DoS Attack Thresholds	
Maximum 1 Minute Connection Rate - low Maximum 1 Minute Connection Rate - high	The number of new unestablished sessions that causes the system to start and stop deleting half-open sessions. Ensure that you enter a lower number in the Low field than you enter in the High field. Possible values are from 1-2147483647 per minute. The default is 400 for low and 500 for high.
Maximum Incomplete Sessions Stop Threshold Maximum Incomplete Sessions Start Threshold	The number of existing half-open sessions that will cause the software to start and stop deleting half-open sessions. Ensure that you enter a lower number in the stop field than you enter in the start field. Possible values are from 1-2147483647. The default is 400 for low and 500 for high.
Thresholds per Host	
Max Sessions Per Host	The number of half-open TCP sessions with the same host destination address that can exist at a time before the software starts deleting half-open sessions to the host. Possible values are 1-4294967295. The default is 50. A large number of half-open sessions can indicate there is a Denial of Service attack against the host.

Table 17-65 Inspection Page (continued)

Element	Description
Max Sessions Blocking Interval (min)	If the maximum sessions per host threshold is reached, the blocking time to apply to help mitigate the potential TCP host-specific denial-of-service (DoS) attack. Possible values are 0-35791 minutes. The default is 0. • If the blocking time value is 0, the software deletes the oldest existing half-open session for the host for every new connection request to the host above the maximum session limit. This ensures that the number of half-open sessions to a given host will never exceed the threshold. • If the blocking time value is greater than 0, the software deletes all existing half-open sessions for the host, then blocks all new connection requests to the host. The software will continue to block all new connection requests until the block-time expires.
Other	
Session Hash Table Size (buckets)	The size of the hash table in terms of buckets. Possible values for the hash table are 1024, 2048, 4096, and 8192. The default is 1024. You should increase the hash table size when the total number of sessions running through the device is approximately twice the current hash size; decrease the hash table size when the total number of sessions is reduced to approximately half the current hash size. Essentially, try to maintain a 1:1 ratio between the number of sessions and the size of the hash table.
Enable Alert Messages	Whether to generate stateful packet inspection alert messages on the console.
Enable Audit Trail Messages	Whether audit trail messages are logged to the syslog server or router.
Permit DHCP Passthrough (Transparent Firewall)	Whether to permit a transparent firewall to forward DHCP packets across the bridge without inspection. Permitting DHCP passthrough overrides an ACL for DHCP packets, so DHCP packets are forwarded even if the ACL is configured to deny all IP packets. Thus, clients on one side of the bridge can get an IP address from a DHCP server on the opposite side of the bridge.
Block Non-SYN Packets	Whether to drop TCP packets that do not belong to an established session. These are TCP packets that do not initiate sessions, that is, the SYN bit is not set in them.
Log Dropped Packets	Whether to create log messages for dropped packets to specify the reason for dropping them.

Related Topics

- Understanding Inspection Rules, page 17-2
- Configuring Inspection Rules, page 17-5
- Using Inspection To Prevent Denial of Service (DoS) Attacks on IOS Devices, page 17-5

Managing Firewall Web Filter Rules

Web filter rules policies define policies for allowing or preventing web traffic based on the requested URL or the applet content of the traffic. For ASA, PIX, and FWSM devices, you can also filter FTP and HTTPS traffic.

How you configure web filter rules is different depending on whether the device uses ASA, PIX or FWSM software as opposed to Cisco IOS Software.

The following topics help you work with web filter rules:

- Understanding Web Filter Rules, page 18-1
- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Configuring Web Filter Rules for IOS Devices, page 18-10
- Configuring Settings for Web Filter Servers, page 18-15

Understanding Web Filter Rules

Web filter rules policies define policies for allowing or preventing web traffic based on the requested URL or the applet content of the traffic. For ASA, PIX, and FWSM devices, you can also filter FTP and HTTPS traffic.

Web, or URL, filtering allows you to control which web sites and web content your users have access to. For example, you might consider some types of content to create a hostile work environment for the people in your organization (for example, web sites that provide pornography). You might consider some web sites to be unsafe and a source of potential viral applications. Using web filter rules, you can block access to these objectionable or unsafe sites.

To filter web requests, you should install an external web filtering server, either Websense or SmartFilter (N2H2). For ASA, PIX, and FWSM devices, these external servers are required for URL, FTP, or HTTPS filtering. For IOS devices, you can also use these servers, but additionally you can create local lists of allowed (always allowed) or blocked (always denied) URLs. You configure the filtering servers in the web filter settings policy; see Configuring Settings for Web Filter Servers, page 18-15.

Tip For IOS devices, you have the option of configuring web filtering using zone-based firewall rules instead of web filter rules, which allows you the additional option of using Trend Micro web filtering servers. For more information, see Chapter 21, “Managing Zone-based Firewall Rules”.

Beside filtering requests based on URL, you can do some applet filtering, stripping out ActiveX or Java applets. You might want to do this to prevent applet downloads from sites you otherwise want to allow if you do not fully trust the site. You can configure your rules to block these applets from specific sites while allowing them from trusted sites.

The policies and procedures for configuring web filter rules differs based on the device type. See the following topics for more information:

- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Configuring Web Filter Rules for IOS Devices, page 18-10

Configuring Web Filter Rules for ASA, PIX, and FWSM Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

Web filter rules policies for ASA, PIX, and FWSM devices define how you want to handle HTTP, FTP, and HTTPS traffic. You can also filter ActiveX and Java applets. Web filter rules permit or deny traffic based on the Universal Resource Locator (URL) address in the web request. If you allow HTTP traffic in your access rules, you can subsequently deny (or drop) traffic if it is directed at an objectionable web or FTP site, or you can strip out ActiveX or Java applets from untrusted sources.

To configure web filtering rules for ASA, PIX, and FWSM devices:

1. Configure the rules that identify traffic that should be subject to filtering, and the traffic that should be exempt from filtering rules (see below for the procedure).
2. Configure web filter settings to identify the URL filtering server and other settings. For more information, see Configuring Settings for Web Filter Servers, page 18-15.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Using Sections to Organize Rules Tables, page 12-20
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19
- Understanding Networks/Hosts Objects, page 6-80
- Understanding and Specifying Services and Service and Port List Objects, page 6-100

Step 1 Do one of the following to open the Web Filter Rules Page (ASA/PIX/FWSM), page 18-3:

- Device view—Select **Firewall > Web Filter Rules** from the Policy selector.
- Policy view—Select **Firewall > Web Filter Rules (PIX/FWSM/ASA)** from the Policy Type select. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes, page 18-5.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10.

Step 3 Configure the rule. Following are the highlights of what you typically need to decide. For specific information on configuring the fields, see Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes, page 18-5.

- Filtering and Type—Whether you are creating a rule that identifies traffic to be filtered (Filter) or exempted from an existing filter rule (Filter Except), and the type of filtering to be done:
  - URL—To filter traffic based on web address.
  - HTTPS—To filter web traffic to secure sites. This does not include SSL VPN traffic.
  - FTP—To filter FTP traffic.
  - ActiveX or Java—To remove ActiveX or Java applets. These options delete all entities within applet or object tags, so you might remove more than just ActiveX or Java applets.
- Source and Destination addresses—If the rule should apply no matter which addresses generated the traffic or their destinations, use "any" as the source or destination. If the rule is specific to a host or network, enter the addresses or network/host objects. For information on the accepted address formats, see Specifying IP Addresses During Policy Definition, page 6-87.
- Service—Primarily defines the port that should be monitored. You must specify some type of TCP service. Typically, you would use the pre-defined services HTTP, HTTPS, or FTP, which should be the same as the type of filtering you are performing, but you can specify any TCP port on your network that might contain the traffic to be filtered.
- Options—The options you want to include, if any. The main options of interest are whether you want to allow traffic if the filtering servers are unavailable, and whether you want to truncate long URLs or URLs that have parameters. Truncating URLs that have parameters is typically a good idea, because if you are going to drop a URL, it is not normally because of a parameter value.

Click **OK** when you are finished defining your rule.

Step 4 If you did not select the desired row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. Order is not as important for web filtering rules, however, because filter except rules always create exceptions to the related filter rule, whether they come before or after the filter rule. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.

Web Filter Rules Page (ASA/PIX/FWSM)

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

Use the Web Filter Rules page for ASA, PIX, and FWSM devices to configure web, or URL, filtering rules. Web filtering is a type of HTTP inspection. If your access rules allow HTTP traffic, you can configure rules to apply server-based web filtering to prevent users from accessing undesirable web servers.

When you configure web filter rules, also configure web filter settings in the **Firewall > Settings > Web Filter** policy. The settings identify the web filtering server and contain other settings that control the overall functioning of the policy. You must configure a web filtering server for your URL, FTP, or HTTPS filter rules to be deployed. For more information, see Web Filter Settings Page, page 18-16.

Tip Rules cannot overlap. For example, if you create two rules with the same, or overlapping, source, destination, and service, you cannot deploy them. Also, you should order any filter-except rules below the filter rule to which they are creating an exemption.

Navigation Path

To access the Web Filter Rules page for ASA, PIX, and FWSM devices, do one of the following:

- (Device view) Select an ASA, PIX, or FWSM device, then select **Firewall > Web Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Web Filter Rules (PIX/FWSM/ASA)** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click an ASA, PIX, or FWSM device and select **Edit Firewall Policies > Web Filter Rules**.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Configuring Settings for Web Filter Servers, page 18-15
- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Using Sections to Organize Rules Tables, page 12-20
- Enabling and Disabling Rules, page 12-20
- Moving Rules and the Importance of Rule Order, page 12-19
- Filtering Tables, page 1-48

Field Reference

Table 18-1 Web Filter Rules Page (ASA, PIX, FWSM)

Element	Description
No.	The ordered rule number.
Source Destination	The source and destination addresses for the rule. The “any” address does not restrict the rule to specific hosts, networks, or interfaces. These addresses are IP addresses for hosts or networks, network/host objects, interfaces, or interface roles. Multiple entries are displayed as separate subfields within the table cell. See Understanding Networks/Hosts Objects, page 6-80.
Service	The services or service objects that specify the protocol and port of the traffic to which the rule applies. Multiple entries are displayed as separate subfields within the table cell. See Understanding and Specifying Services and Service and Port List Objects, page 6-100.

Table 18-1 Web Filter Rules Page (ASA, PIX, FWSM) (continued)

Element	Description
Type	The type of filtering action for the rule, either filtering the identified traffic, or exempting the identified traffic from filtering (Filter Except). For a full explanation, see Edit Web Filter Type Dialog Box, page 18-8.
Options	Additional configuration options for the selected protocol, if any. For detailed descriptions, see Edit Web Filter Options Dialog Box, page 18-9.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of the rule, if any.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Query	Click this button to run a policy query, which can help you evaluate your rules and identify ineffective rules. See Generating Policy Query Reports, page 12-28
Find and Replace button (binoculars icon)	Click this button to search for various types of items within the table and to optionally replace them. See Finding and Replacing Items in Rules Tables, page 12-16.
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope or section. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Add Row button	Click this button to add a rule to the table after the selected row using the Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes, page 18-5. If you do not select a row, the rule is added at the end of the local scope. For more information about adding rules, see Adding and Removing Rules, page 12-9.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells. For more information, see Editing Rules, page 12-10.
Delete Row button	Click this button to delete the selected rule.

Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

Use the Add and Edit PIX/ASA/FWSM Web Filter Rule dialog boxes to configuring web filtering rules for these types of devices.

Navigation Path

From the Web Filter Rules Page (ASA/PIX/FWSM), page 18-3, click the **Add Row** button or select a row and click the **Edit Row** button.

Related Topics

- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Understanding Web Filter Rules, page 18-1
- Configuring Settings for Web Filter Servers, page 18-15

Field Reference

Table 18-2 Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes

Element	Description
Enable Rule	Whether to enable the rule, which means the rule becomes active when you deploy the configuration to the device. Disabled rules are shown overlain with hash marks in the rule table. For more information, see Enabling and Disabling Rules, page 12-20.
Filtering	The type of rule you are defining: • Filter—The rule filters the identified type of traffic between source and destination. • Filter Except—The rule creates an exemption to a filter rule. The identified traffic between the source and destination is not filtered.
Type	The type of traffic that should be filtered (or exempted from filtering) for this rule. For filtering that uses an external server, consult the documentation for your version of the server to determine if it supports that type of filtering. Configure the filtering server on the Web Filter Settings Page, page 18-16. • URL—HTTP traffic. Filtering is done using an external filtering server. • HTTPS—HTTPS traffic. This does not include traffic associated with an SSL VPN. Filtering is done using an external filtering server. • Java—Remove Java applets from HTTP traffic if they are identified on applet tags. The rule does not remove Java applets from SSL VPN traffic. If the applet tag spans packets, or the code in the tags is larger than the MTU, the Java applet is not removed. • ActiveX—Remove ActiveX or Java applets from HTTP traffic. The rule removes any item within object or applet tags, which might also remove images and multimedia objects. The rule might not remove applets from SSL VPN traffic. If the object tag spans packets, or the code in the tags is larger than the MTU, the object is not removed. • FTP—FTP traffic. Filtering is done using an external filtering server.

Table 18-2 Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes (continued)

Element	Description
Sources Destinations	The source or destination of the traffic. You can enter more than one value by separating the items with commas. You can enter any combination of the following address types to define the source or destination of the traffic. For more information, see Specifying IP Addresses During Policy Definition, page 6-87. • Network/host object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list. • Host IP address, for example, 10.10.10.100. • Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0. • A range of IP addresses, for example, 10.10.10.100-10.10.10.200. • An IP address pattern in the format 10.10.0.10/255.255.0.255, where the mask is a discontiguous bit mask (see Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81).
Services	The services that define the port number of the traffic to act on. You can enter more than one value by separating the items with commas. The service must use TCP. Your specification defines the port that you want filtered (the service name has no meaning). For example, if you want to filter port 80, use the HTTP service object. If HTTP traffic on your network uses a different port, specify TCP/port number (for example, TCP/8080). You can enter TCP by itself to filter all ports. You can enter any combination of service objects and service types (which are typically a protocol and port combination). If you type in a service, you are prompted as you type with valid values. You can select a value from the list and press Enter or Tab. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.
Allow traffic if URL Filter Server unavailable (URL, FTP, HTTPS only)	Whether to permit unfiltered traffic on outbound connections if all of the URL filtering servers are unavailable. If you do not select this option, all affected outbound traffic (HTTP, FTP, or HTTPS) is blocked until at least one filtering server becomes available.
Block connection to HTTP Proxy Server (URL only)	Whether to prevent users from connecting to an HTTP proxy server.
Truncate CGI request by removing CGI parameters (URL only)	When a URL has a parameter list starting with a question mark (?), such as a CGI script, whether to truncate the URL sent to the filtering server by removing all characters after and including the question mark.
Block outbound requests if absolute FTP path is not provided (FTP only)	Whether to prevent interactive FTP sessions that do not provide the entire directory path when the user tries to change directories.

Table 18-2 Add and Edit PIX/ASA/FWSM Web Filter Rule Dialog Boxes (continued)

Element	Description
Long URL (URL only)	How to handle URLs that are longer than the maximum allowed by the filtering server: 4 KB for Websense, 3 KB for Smartfilter (N2H2). Many times, long URLs are due to parameter lists, and you can use the **Truncate CGI request by removing CGI parameters** option to handle those URLs. For other long URLs, select from the following options: • Drop—Drop the long URL request. • Truncate—Truncate the URL request to only the hostname or IP address portion of the URL. • Deny—Deny the URL request.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	An optional description of the rule (up to 1024 characters).

Edit Web Filter Type Dialog Box

Use the Edit Web Filter Type dialog box to edit the type of filtering to be done by a web filter rule for ASA, PIX, and FWSM devices.

Navigation Path

Right-click the Type cell in a web filter rule for ASA/PIX/FWSM (on the Web Filter Rules Page (ASA/PIX/FWSM), page 18-3) and select **Edit Web Filter Type**. You can edit the type for one row at a time.

Field Reference

Table 18-3 Edit Web Filter Type Dialog Box

Element	Description
Filtering	The type of rule you are defining: • Filter—The rule filters the identified type of traffic between source and destination. • Filter Except—The rule creates an exemption to a filter rule. The identified traffic between the source and destination is not filtered.

Table 18-3 Edit Web Filter Type Dialog Box (continued)

Element	Description
Type	The type of traffic that should be filtered (or exempted from filtering) for this rule. For filtering that uses an external server, consult the documentation for your version of the server to determine if it supports that type of filtering. Configure the filtering server on the Web Filter Settings Page, page 18-16. • URL—HTTP traffic. Filtering is done using an external filtering server. • HTTPS—HTTPS traffic. This does not include traffic associated with an SSL VPN. Filtering is done using an external filtering server. • Java—Remove Java applets from HTTP traffic if they are identified on applet tags. The rule does not remove Java applets from SSL VPN traffic. If the applet tag spans packets, or the code in the tags is larger than the MTU, the Java applet is not removed. • ActiveX—Remove ActiveX or Java applets from HTTP traffic. The rule removes any item within object or applet tags, which might also remove images and multimedia objects. The rule might not remove applets from SSL VPN traffic. If the object tag spans packets, or the code in the tags is larger than the MTU, the object is not removed. • FTP—FTP traffic. Filtering is done using an external filtering server.

Edit Web Filter Options Dialog Box

Use the Edit Web Filter Options dialog box to edit the filtering options defined for a web filter rule for ASA, PIX, and FWSM devices.

The options displayed on this dialog box differ depending on the type of filtering configured for the rule. For some types, there are no options and the dialog box is empty. The reference table below includes all possible options.

Navigation Path

Right-click the Options cell in a web filter rule for ASA/PIX/FWSM (on the Web Filter Rules Page (ASA/PIX/FWSM), page 18-3) and select **Edit Web Filter Type**. You can edit the type for one row at a time.

Field Reference

Table 18-4 Edit Web Filter Options Dialog Box

Element	Description
Allow traffic if URL Filter Server unavailable (URL, FTP, HTTPS only)	Whether to permit unfiltered traffic on outbound connections if all of the URL filtering servers are unavailable. If you do not select this option, all affected outbound traffic (HTTP, FTP, or HTTPS) is blocked until at least one filtering server becomes available.

Table 18-4 Edit Web Filter Options Dialog Box (continued)

Element	Description
Block connection to HTTP Proxy Server (URL only)	Whether to prevent users from connecting to an HTTP proxy server.
Truncate CGI request by removing CGI parameters (URL only)	When a URL has a parameter list starting with a question mark (?), such as a CGI script, whether to truncate the URL sent to the filtering server by removing all characters after and including the question mark.
Block outbound requests if absolute FTP path is not provided (FTP only)	Whether to prevent interactive FTP sessions that do not provide the entire directory path when the user tries to change directories.
Long URL (URL only)	How to handle URLs that are longer than the maximum allowed by the filtering server: 4 KB for Websense, 3 KB for Smartfilter (N2H2). Many times, long URLs are due to parameter lists, and you can use the **Truncate CGI request by removing CGI parameters** option to handle those URLs. For other long URLs, select from the following options: • Drop—Drop the long URL request. • Truncate—Truncate the URL request to only the hostname or IP address portion of the URL. • Deny—Deny the URL request.

Configuring Web Filter Rules for IOS Devices

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Web filter rules policies for IOS devices define how you want to handle HTTP traffic. The web filter rules are a type of inspection rule that permits or denies traffic based on the Universal Resource Locator (URL) address in the web request. If you allow HTTP traffic on an interface in your access rules, you can subsequently deny (or drop) traffic if it is directed at an objectionable web site.

To configure web filtering rules for IOS devices:

1. Configure the interfaces that should filter web traffic (see below for the procedure).
2. Configure the local web filtering list to identify web sites that should always be permitted or denied (see below for the procedure).
3. Configure web filter settings to identify the URL filtering server and other settings. For more information, see Configuring Settings for Web Filter Servers, page 18-15.

Tip You can also configure web filtering as a zone based firewall rule. For more information, see Adding Zone-Based Firewall Rules, page 21-13.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Understanding Interface Role Objects, page 6-73
- Understanding Networks/Hosts Objects, page 6-80

Step 1 Do one of the following to open the Web Filter Rules Page (IOS), page 18-12:

- Device view—Select **Firewall > Web Filter Rules** from the Policy selector.
- Policy view—Select **Firewall > Web Filter Rules (IOS)** from the Policy Type select. Select an existing policy or create a new one.

Step 2 Configure the interfaces on which you will filter HTTP traffic. Create rules for each interface on which you will enable filtering:

a. Select the Web Filter Rules tab if it is not already selected and do one of the following to open the IOS Web Filter Rule and Applet Scanner Dialog Box, page 18-13:

- To create a new rule, right-click inside the work area and select **Add Row**.
- To edit an existing rule, right-click the rule and select **Edit Row**.

b. Identify the interface for which this rule applies. You can either enter the interface name or click **Select** to select it or an interface role from the list. Also configure the following:

- Traffic direction with respect to the interface—Typically, you want to select **In** so that undesired traffic is dropped before the device spends more time processing the packet.
- Java applet scanning—If you enable web filtering on an interface, Java applets are inspected, which can affect performance. Typically, you want to enable Java applet scanning so that you can identify permitted and denied sources and avoid the scanning of denied applets. If you want to configure both permitted and denied sources for an interface, you must configure two rules for the interface.

c. Click **OK** to add the rule to the web filtering rules table.

Step 3 (Optional) Configure the list of exclusive domains, which define the local filtering list. This list is applied before web requests are sent to the external web filtering server (defined on the Web Filter Settings Page, page 18-16). If you know there are web sites that you will always permit (such as your organization's web site) or deny, configure them in the local list. Configure as many rules as needed to define the complete list.

a. Click the **Exclusive Domains** tab and do one of the following to open the IOS Web Filter Exclusive Domain Name Dialog Box, page 18-14.

- To create a new rule, right-click inside the work area and select **Add Row**.
- To edit an existing rule, right-click the rule and select **Edit Row**.

b. Select whether you are permitting or denying the specified domains, and enter the domain names or host IP addresses. You can enter either full domain names (the names of specific web sites) or partial names (for entire domains you want to treat the same way).

c. Click **OK** to add your exclusive domain rule to the policy.

Web Filter Rules Page (IOS)

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the Web Filter Rules page for IOS devices to configure web, or URL, filtering rules. Web filtering is a type of HTTP inspection. If your access rules allow HTTP traffic on an interface, you can configure rules to apply local and server-based web filtering to prevent users from accessing undesirable web servers.

When you configure web filter rules, also configure web filter settings in the **Firewall > Settings > Web Filter** policy. The settings identify the web filtering server and contain other settings that control the overall functioning of the policy. For example, you can use the settings policy to allow all web traffic if the filtering server becomes unavailable. For more information, see Web Filter Settings Page, page 18-16.

Tip You can also configure web filtering as a zone based firewall rule. For more information, see Zone-based Firewall Rules Page, page 21-58.

Navigation Path

To access the Web Filter Rules page for IOS devices, do one of the following:

- (Device view) Select an IOS device and select **Firewall > Web Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Web Filter Rules (IOS)** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click an IOS device and select **Edit Firewall Policies > Web Filter Rules**.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Configuring Web Filter Rules for IOS Devices, page 18-10
- Chapter 18, “Managing Firewall Web Filter Rules”

Field Reference

Table 18-5 Web Filter Rules Page (IOS)

Element	Description
Web Filter Rules tab	The URL filtering rules defined for the policy. Each rule shows the interface on which it is defined, whether the rule is applied to incoming or outgoing traffic, and the permitted or denied Java applet sources if Java applet scanning is enabled. You might have more than one rule for an interface if you configure both a permit and deny list for Java applet scanning. • To add a rule, click the Add Row button and fill in the IOS Web Filter Rule and Applet Scanner Dialog Box, page 18-13. • To edit a rule, select it and click the Edit Row button. • To delete a rule, select it and click the Delete Row button.
Exclusive Domains tab	The local web filter list. This list is checked before web requests are sent to the filtering server and applies to all interfaces on which you configure web filtering. If you know there are specific domains that you will always allow (such as your organization's own domain name), or disallow, you can list them here. By configuring a local filter list, you can improve performance because the device does not need to wait for a response from the filtering server. • To add a domain, click the Add Row button and fill in the IOS Web Filter Exclusive Domain Name Dialog Box, page 18-14. • To edit a domain, select it and click the Edit Row button. • To delete a domain, select it and click the Delete Row button.

IOS Web Filter Rule and Applet Scanner Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the IOS Web Filter Rule and Applet Scanner dialog box to create web filtering rules for IOS devices.

Navigation Path

To open this dialog box, select the Web Filter Rules tab on the Web Filter Rules Page (IOS), page 18-12, click **Add Row** to create a new rule, or select a row and click **Edit Row** to edit an existing rule.

Related Topics

- Configuring Web Filter Rules for IOS Devices, page 18-10
- Understanding Web Filter Rules, page 18-1
- Chapter 18, "Managing Firewall Web Filter Rules"
- Configuring Settings for Web Filter Servers, page 18-15

Field Reference

Table 18-6 IOS Web Filter Rule and Applet Scanner Dialog Box

Element	Description
Enable Web Filtering	Whether to enable the web filtering rule.
Interface	The interface or interface role to which the rule is assigned. Enter the name of the interface or the interface role, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. See Understanding Interface Role Objects, page 6-73.
Traffic Direction	The direction of the traffic to which this rule applies: • In—Packets entering an interface. • Out—Packets exiting an interface.
Java Applet Scanning Enable Java Applet Scanner	If you select **Enable Java Applet Scanning**, the device checks for the presence of Java applets in HTTP traffic coming from web servers to internal hosts. If a Java applet is present and the web server (applet source) is in the list of permitted sources, the Java applet is left unmodified in the HTTP traffic. Otherwise, the Java applets are removed from HTTP pages. **Tip** When you enable web filtering, Java applets are inspected, which can affect performance. By enabling the Java applet scanner, you can identify a list of permitted or denied sources and avoid inspection for those applets. Even if you do not want to deny any sources, enable scanning and permit the **any** source.
Permit Traffic Applet Sources	The list of permitted or denied source addresses for Java applets. To configure a list of permitted or denied sources: • Select either **Permit from Specified Sources** or **Deny from Specified Sources**. If you want to create both a permit and deny list, create two separate web filter rules. If you do not configure a permit list, all sources are denied. • Enter the list of permitted or denied addresses in the **Applet Sources** field. The list can include host IP addresses, network addresses, address ranges, or network/host objects, but cannot include domain names. Separate multiple addresses with commas. For more information on entering addresses, see Specifying IP Addresses During Policy Definition, page 6-87.

IOS Web Filter Exclusive Domain Name Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the IOS Web Filter Exclusive Domain Name dialog box to configure local web filtering rules for IOS devices. You can create a list of permitted or denied domain names or IP addresses. The device checks this list before forwarding web requests to your web filtering server.

Using local filtering saves the wait time for getting a response from the server when a user requests a web site that you know you will either always permit or always deny.

Navigation Path

To open this dialog box, select the Exclusive Domains tab on the Web Filter Rules Page (IOS), page 18-12, click **Add Row** to create a new rule, or select a row and click **Edit Row** to edit an existing rule.

Related Topics

- Configuring Web Filter Rules for IOS Devices, page 18-10
- Understanding Web Filter Rules, page 18-1
- Chapter 18, “Managing Firewall Web Filter Rules”

Field Reference

Table 18-7 IOS Web Filter Exclusive Domain Name Dialog Box

Element	Description
Traffic	Whether you want to permit access to the listed web sites or deny access to them.
Domain Name	The domain names or host IP addresses of web sites that you are permitting or denying. Separate multiple entries with commas. For domain names, you can enter a full or partial name. For example, cisco.com covers all web servers on the cisco.com domain, whereas www.cisco.com specifies only the www web server.

Configuring Settings for Web Filter Servers

Use the Web Filter settings policy to configure the web filter servers and other settings to use with your web filter rules policy. You can use Websense or Smartfilter (N2H2) filtering servers, or no external servers (for IOS devices).

You must install and configure the web filter servers as directed by the documentation for the server before configuring and deploying this policy. Security Manager cannot confirm that the servers exist or that they are configured correctly.

Tip These settings work only with the web filter rules policy. The web servers you configure here are not used with zone based firewall rules policies that configure web content filtering.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Configuring Web Filter Rules for IOS Devices, page 18-10

Step 1 Do one of the following to open the Web Filter Settings Page, page 18-16:

- (Device view) Select **Firewall > Settings > Web Filter** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Web Filter** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the type of web filtering server you use in the **Web Filter Server Type** field, and then add the servers to the table of web filtering servers. If you have more than one server, add them in priority order; the first server in the list is the primary server.

- To add a server, click the Add Row button and fill in the Web Filter Server Configuration Dialog Box, page 18-19.
- To edit a server, select it and click the Edit Row button.
- To delete a server, select it and click the Delete Row button.

Step 3 The bottom half of the settings policy includes device-specific options that you can also configure. For specific information on each setting, see Web Filter Settings Page, page 18-16. The following is an overview of the settings:

- IOS devices—The most interesting setting is **Allow Traffic when Servers Unreachable**, which determines whether you allow any web connections if the filtering servers are unavailable. If you do not select this option, all web traffic is cut off if the servers go offline for any reason.

  The remaining settings configure logging and cache size options.
- ASA, PIX, FWSM devices—These options configure the cache size and buffer limits used with the filtering servers. You can also control whether the cached responses include both source and destination (if you have different filtering policies per user) or destination only (one policy for all), as configured in the filtering server.

Web Filter Settings Page

Use the Web Filter settings page to configure the web filter servers and other settings to use with your web filter rules policy.

You must install and configure the web filter servers as directed by the documentation for the server before configuring and deploying this policy. Security Manager cannot confirm that the servers exist or that they are configured correctly.

Tip These settings work only with the web filter rules policy. The web servers you configure here are not used with zone based firewall rules policies that configure web content filtering.

Navigation Path

To access the Web Filter settings page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Web Filter** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Web Filter** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Settings > Web Filter**.

Related Topics

- Understanding Web Filter Rules, page 18-1
- Configuring Settings for Web Filter Servers, page 18-15
- Configuring Web Filter Rules for ASA, PIX, and FWSM Devices, page 18-2
- Configuring Web Filter Rules for IOS Devices, page 18-10

Field Reference

Table 18-8 Web Filter Page

Element	Description
Web Filter Server Type	The type of web filter server you are using: • None—You are not using web filter servers. • Websense—You use Websense servers. • Secure Computing SmartFilter/N2H2—You use Smartfilter servers. If you select this option, you can specify the server port to use for communication in the **Port** field. **Tip** If you change this setting, you are prompted to remove the existing list of servers from the table. Clicking **Yes** does not clear the table. The prompt is to remind you that the list might contain the wrong type of servers.
Web Filter Servers table	The servers that the device should use for web filtering. Enter the servers in priority order; the device uses the first one in the list until it fails to respond, and moves to the next server in the list until it gets a response. If you select None for filter type, this list is ignored. • To add a server, click the **Add Row** button and fill in the Web Filter Server Configuration Dialog Box, page 18-19. • To edit a server, select it and click the **Edit Row** button. • To delete a server, select it and click the **Delete Row** button.
IOS Specific Settings	
Allow Traffic when Servers Unreachable	Whether the device should allow web traffic if the web filter servers are not responding. If you do not select this option, all web access is prevented until the servers come back online. If you allow web traffic when the servers are down, the web requests are not filtered and access to all web servers is allowed.
Enable Alerts	Whether to generate stateful packet inspection alert messages on the console.
Enable Audit Trail	Whether audit trail messages are logged to the syslog server or router.
Enable Web Filter Server Logging	Whether to send system messages to the URL filtering server for logging. The device sends a log request immediately after the URL lookup request. The log request contains the URL, hostname, source IP address, and the destination IP address. The server records the log request into its own log server so you can view this information as necessary.

Table 18-8 Web Filter Page (continued)

Element	Description
Cache Size	The maximum number of destination IP addresses (and their authorization status) that can be cached in the device. The default value is 5000. When the cache reaches 80% full, the device starts removing older inactive entries.
Maximum Requests	The maximum number of outstanding requests that can exist at any given time. If the specified number is exceeded, new requests are dropped. The default is 1000.
Packet Buffer	The maximum number of HTTP responses that can be stored in the packet buffer of the device while it waits for the web filter server to allow or deny the request. The device drops responses when the maximum is reached. The default (and maximum) value is 200. When users make web requests, the device simultaneously sends the request to the web site and to the web filtering server. If the response from the web site is received before the server provides a permit or deny response, the device keeps the request in the packet buffer until it gets a response from the server. The response is removed from the buffer when the server responds or if the device determines that the server is unavailable and you also selected **Allow Traffic when Servers Unreachable**.
PIX/ASA/FWSM Specific Settings	
Cache Match Criteria	How to cache web requests: • Source and Destination—Cache entries are based on both the address initiating the request and the destination web address. Select this mode if users do not share the same filtering policy on the filtering server. • Destination—Cache entries are based on the destination web address. Select this mode if all users share the same filtering policy on the filtering server.
URL Buffer Memory (ASA 7.2+, PIX 7.2+ only.)	The size of the URL buffer memory pool in KB. Values are 2 to 10240.
Maximum Allowed URL Size (ASA 7.2+, PIX 7.2+ only.)	The maximum allowed URL size in KB for each URL being buffered. The possible values differ depending on server type: • Websense—From 2 to 4. • Smartfilter (N2H2)—2 or 3.
Cache Size	The size of the cache, in KB, for storing responses from the filtering server. Values are 1 to 128. Caching stores URL access privileges in memory on the security appliance. When a host requests a connection, the security appliance first looks in the URL cache for matching access privileges instead of forwarding the request to the Websense server.

Table 18-8 Web Filter Page (continued)

Element	Description
URL Block Buffer Limit	The size of the buffer for storing web server responses while waiting for a filtering decision from the filtering server. The values are 1 to 128, which specifies the number of 1550-byte blocks.

Web Filter Server Configuration Dialog Box

Use the Web Filter Server Configuration dialog box to configure the external web filter servers you want to use with your Web Filter Rules policies. You can configure Websense or Smartfilter (N2H2) servers.

Navigation Path

From the Web Filter Settings Page, page 18-16, click **Add Row** beneath the Web Filter Servers table, or select a row and click **Edit Row**.

Related Topics

- Configuring Settings for Web Filter Servers, page 18-15
- Understanding Web Filter Rules, page 18-1
- Chapter 18, "Managing Firewall Web Filter Rules"

Field Reference

Table 18-9 Web Filter Server Configuration Dialog Box

Element	Description
Common	
IP Address	The IP address of the web filter server.
Timeout	The length of time, in seconds, that the device will wait for a response from the web filter server. The default is 5 seconds. If the request times out, the device tries the next server, if you configure more than one.
PIX/ASA/FWSM Specific Settings	
Interface	The network interface where the authentication server resides, for example, FastEthernet0. If not specified, the default is inside. Enter the name of the interface or the interface role that identifies it, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list.
Protocol	The protocol to use when communicating with the web filtering server. Select the option for which the server is configured: • TCP (version 1) • TCP version 4 • UDP version 4
Connection Number	(Optional) The maximum number of TCP connections allowed between the device and the server.

Table 18-9 Web Filter Server Configuration Dialog Box (continued)

Element	Description
IOS Specific Settings	
Retransmit	The number of times the device will retransmit a request when the server does not respond. The default value is two times.
Port	The port number that the server listens on. The default port is 15868.

Managing Firewall Botnet Traffic Filter Rules

Malware is malicious software that is installed on an unknowing host. Malware that attempts network activity such as sending private data (passwords, credit card numbers, key strokes, or proprietary data) can be detected by the Botnet Traffic Filter when the malware starts a connection to a known bad IP address. The Botnet Traffic Filter checks incoming and outgoing connections against a dynamic database of known bad domain names and IP addresses, and then logs any suspicious activity.

You can also supplement the Cisco dynamic database with addresses on the block list, of your choosing by adding them to a static block list; if the dynamic database includes blocked addresses that you think should not be blocked, you can manually enter them into a static allow list. Addresses on the allow list still generate syslog messages, but because you are only targeting blocked syslog messages, they are informational. If you do not want to use the Cisco dynamic database at all, because of internal requirements, you can use the static block list alone if you can identify all the malware sites that you want to target.

Related Topics

- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9

Understanding Botnet Traffic Filtering

Botnet Traffic Filter Address Categories

Addresses monitored by the Botnet Traffic Filter include:

- **Known malware addresses**—These addresses are on the block list identified by the dynamic database and the static block list.
- **Known allowed addresses**—These addresses are on the allow list. To be allowed, an address must be blocked by the dynamic database and also identified by the static allow list.
- **Ambiguous addresses**—These addresses are associated with multiple domain names, but not all of these domain names are on the block list. These addresses are on the graylist.
- **Unlisted addresses**—These addresses are unknown, and not included on any list.

Botnet Traffic Filter Actions for Known Addresses

You can configure the Botnet Traffic Filter to log suspicious activity, and you can optionally configure it to block suspicious traffic automatically.

Unlisted addresses do not generate any syslog messages, but addresses on the block list, allow list, and graylist generate syslog messages differentiated by type.

Botnet Traffic Filter Databases

The Botnet Traffic Filter uses two databases for known addresses. You can use both databases together, or you can disable use of the dynamic database and use the static database alone. This section includes the following topics:

- Information About the Dynamic Database
- Information About the Static Database

Information About the Dynamic Database

The Botnet Traffic Filter can receive periodic updates for the dynamic database from the Cisco update server. This database lists thousands of known bad domain names and IP addresses.

The security appliance uses the dynamic database as follows:

1. When the domain name in a DNS reply matches a name in the dynamic database, the Botnet Traffic Filter adds the name and IP address to the DNS reverse lookup cache.
2. When the infected host starts a connection to the IP address of the malware site, the security appliance sends a syslog message informing you of the suspicious activity.
3. In some cases, the IP address itself is supplied in the dynamic database, and the Botnet Traffic Filter logs any traffic to that IP address without having to inspect DNS requests.

Note To use the database, be sure to configure a domain name server for the security appliance so that it can access the URL.

To use the domain names in the dynamic database, you need to enable DNS packet inspection with Botnet Traffic Filter snooping; the security appliance looks inside the DNS packets for the domain name and associated IP address.

Information About the Static Database

You can manually enter domain names or IP addresses (host or subnet) that you want to tag as bad names in a block list. You can also enter names or IP addresses in an allow list, so that names or addresses that appear on both the allow list and the dynamic block list are identified only as allowed addresses in syslog messages and reports.

You can alternatively enable DNS packet inspection with Botnet Traffic Filter snooping. With DNS snooping, when an infected host sends a DNS request for a name on the static database, the security appliance looks inside the DNS packets for the domain name and associated IP address and adds the name and IP address to the DNS reverse lookup cache.

Related Topics

- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9

Task Flow for Configuring the Botnet Traffic Filter

To configure the Botnet Traffic Filter, follow these steps:

Step 1 Enable use of a DNS server.

This procedure enables security appliance use of a DNS server. In multiple context mode, enable DNS per context.

For more information, see DNS Page, page 52-14

Step 2 Enable use of the dynamic database.

This procedure enables database updates from the Cisco update server, and also enables use of the downloaded dynamic database by the security appliance. Disallowing use of the downloaded database is useful in multiple context mode so you can configure use of the database on a per-context basis.

For more information, see Configuring the Dynamic Database, page 19-4

Step 3 (Optional) Add static entries to the database.

This procedure lets you augment the dynamic database with domain names or IP addresses that you want to block or allow. You might want to use the static database instead of the dynamic database if you do not want to download the dynamic database over the Internet.

For more information, see Adding Entries to the Static Database, page 19-5

Step 4 Enable DNS snooping.

This procedure enables inspection of DNS packets, compares the domain name with those in the dynamic database or the static database (when a DNS server for the security appliance is unavailable), and adds the name and IP address to the DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter logging function when connections are made to the suspicious address.

For more information, see Enabling DNS Snooping, page 19-6

Step 5 Enable traffic classification and actions for the Botnet Traffic Filter.

This procedure enables the Botnet Traffic Filter, which compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message for any matching traffic or drops that traffic.

For more information, see Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6

Step 6 Monitor and Mitigate Botnet Activity.

After configuring the Botnet Traffic Filter on a device, the device will begin generating syslog messages to notify you of botnet activity. You should verify the syslog configuration on the device so that messages are appropriately logged and that notifications are sent as needed. As malicious traffic is identified, you will need to perform necessary actions to stop such traffic and to clean any infected computers that are generating the malicious traffic.

For more information, see the following references:

1. Chapter 54, “Configuring Logging Policies on Firewall Devices”
2. Monitoring and Mitigating Botnet Activity, page 69-61
3. Understanding Firewall Summary Botnet Reports, page 70-15

Configuring the Dynamic Database

This procedure enables database updates, and also enables use of the downloaded dynamic database by the security appliance.

In multiple context mode, you enable downloading of the dynamic database on the System context so that it is available to all security contexts. You can then decide, on a per-context basis, whether to enable use of the dynamic database or not.

By default, downloading and using the dynamic database is disabled.

Related Topics

- Dynamic Blacklist Configuration Tab, page 19-10
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Adding Entries to the Static Database, page 19-5
- Enabling DNS Snooping, page 19-6
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Botnet Traffic Filter Rules Page, page 19-9

Before You Begin

Enable security appliance use of a DNS server (see DNS Page, page 52-14). In multiple context mode, enable DNS per context.

Step 1 Do one of the following:

- (Device view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy Type selector. Select an existing policy or create a new one.

Note For devices in multiple context mode, you enable downloading of the dynamic database on the System context and enable use of the dynamic database on each security context, as needed.

This opens the Botnet Traffic Filter Rules Page, page 19-9.

Step 2 On the Dynamic Blacklist Configuration tab, select **Enable Dynamic Blacklist From Server** to enable downloading of the dynamic database.

Note In multiple context mode, you enable downloading of the dynamic database on the System context.

This setting enables downloading of the dynamic database from the Cisco update server. If you do not have a database already installed on the security appliance, it downloads the database after approximately 2 minutes. The update server determines how often the security appliance polls the server for future updates, typically every hour.

Step 3 (Multiple context mode only) Click **Save** to save the changes to the System context. Then change to the context where you want to configure the Botnet Traffic Filter, select **Firewall > Botnet Traffic Filter Rules** for that context, and then proceed to Step 4.

Step 4 On the Dynamic Blacklist Configuration tab, select **Use Dynamic Blacklist** to enable use of the dynamic database.

Note In multiple context mode, these settings are disabled on the System context.

Adding Entries to the Static Database

The static database lets you augment the dynamic database with domain names, IP addresses, or network addresses that you want to block or allow. For more information, see Understanding Botnet Traffic Filtering, page 19-1.

Related Topics

- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Configuring the Dynamic Database, page 19-4
- Enabling DNS Snooping, page 19-6
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Botnet Traffic Filter Rules Page, page 19-9

Before You Begin

- Enable security appliance use of a DNS server (see DNS Page, page 52-14). In multiple context mode, enable DNS per context.

Step 1 Do one of the following:

- (Device view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy Type selector. Select an existing policy or create a new one.

Note For devices in multiple context mode, you configure the static database on the security context.

This opens the Botnet Traffic Filter Rules Page, page 19-9.

Step 2 On the Whitelist / Blacklist tab, click the **Add Rows** button that corresponds with the type of entry you are adding (Whitelist or Blacklist).

This opens the Device Whitelist or Device Blacklist Dialog Box, page 19-15.

Step 3 In the Domain or IP Address field, enter one or more domain names, IP addresses, and IP address/netmasks. Enter multiple entries separated by commas or on separate lines. You can enter up to 1000 entries for each type.

Step 4 Click **OK**.

Enabling DNS Snooping

This procedure enables inspection of DNS packets and enables Botnet Traffic Filter snooping, which compares the domain name with those on the dynamic database or static database, and adds the name and IP address to the Botnet Traffic Filter DNS reverse lookup cache. This cache is then used by the Botnet Traffic Filter logging function when connections are made to the suspicious address.

The default configuration for DNS inspection inspects all UDP DNS traffic on all interfaces, and does not have Botnet Traffic Filter snooping enabled. We suggest that you enable Botnet Traffic Filter snooping only on interfaces where external DNS requests are going. Enabling Botnet Traffic Filter snooping on all UDP DNS traffic, including that going to an internal DNS server, creates unnecessary load on the security appliance.

Note TCP DNS traffic is not supported.

Related Topics

- Configure DNS Dialog Box, page 17-19
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Configuring the Dynamic Database, page 19-4
- Adding Entries to the Static Database, page 19-5
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Botnet Traffic Filter Rules Page, page 19-9

Step 1 You must first configure DNS inspection for traffic that you want to snoop using the Botnet Traffic Filter. See Chapter 17, "Managing Firewall Inspection Rules".

Step 2 While defining a new inspection rule or editing an existing inspection rule, select DNS as the protocol you want to inspect.

The Configure button to the right of the Selected Protocol field becomes active.

Step 3 Click **Configure**.

This opens the Configure DNS Dialog Box, page 17-19.

Step 4 To enable DNS snooping, select **Enable Dynamic Filter Snooping**.

Step 5 Click **OK**.

Enabling Traffic Classification and Actions for the Botnet Traffic Filter

This procedure enables the Botnet Traffic Filter, which compares the source and destination IP address in each initial connection packet to the IP addresses in the dynamic database, static database, DNS reverse lookup cache, and DNS host cache, and sends a syslog message for any matching traffic. The

Botnet Traffic Filter can also drop the connection when matching traffic is encountered. For a particular interface, you can specify only one enable rule that identifies the traffic that is subject to Botnet Traffic Filtering; however, you can specify multiple drop rules to identify traffic that should be dropped by the Botnet Traffic Filter.

The DNS snooping is enabled separately (see Enabling DNS Snooping, page 19-6). Typically, for maximum use of the Botnet Traffic Filter, you need to enable DNS snooping, but you can use Botnet Traffic Filter logging independently if desired. Without DNS snooping for the dynamic database, the Botnet Traffic Filter uses only the static database entries, plus any IP addresses in the dynamic database; domain names in the dynamic database are not used.

What You Need To Know About Botnet Traffic Classification ACLs

When you configure the enable and drop rules, you have the option of specifying an extended ACL policy object to limit the traffic to which Botnet Traffic Filtering will be applied. If you do not specify an ACL object, filtering is done for all traffic: this is equivalent to specifying an ACL with the single rule permit IP any any.

If you want to specify an ACL so that filtering is performed on less than all traffic, keep the following in mind:

- Permit rules identify the traffic that is subject to Botnet Traffic Filtering. In drop rules, permit entries identify the traffic that the ASA is allowed to drop.
- Deny rules identify the traffic that should not be subject to filtering. The Botnet Traffic Filter ignores traffic that matches deny entries.
- The ACL that you select for drop rules should be a subset of the ACL used in the enable rules for the interface. For traffic to be dropped, there must not only be a permit rule in the drop rule's ACL, the traffic must also fall under a permit rule in the enable rule's ACL. This is because the drop rule is not considered until traffic permitted in an enable rule has first been identified as blocked.

We recommend enabling the Botnet Traffic Filter on all traffic on the Internet-facing interface, and enabling dropping of traffic with a severity of moderate and higher.

Related Topics

- Traffic Classification Tab, page 19-11
- BTF Enable Rules Editor, page 19-12
- BTF Drop Rules Editor, page 19-13
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Configuring the Dynamic Database, page 19-4
- Adding Entries to the Static Database, page 19-5
- Enabling DNS Snooping, page 19-6
- Botnet Traffic Filter Rules Page, page 19-9

Step 1 Do one of the following:

- (Device view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy Type selector. Select an existing policy or create a new one.

Note For devices in multiple context mode, you configure traffic classification on the security context.

This opens the Botnet Traffic Filter Rules Page, page 19-9.

Step 2 To enable the Botnet Traffic Filter on specified traffic, follow these steps:

a. On the Traffic Classification tab, click **Add Row** under the Enable Rules table.

This opens the BTF Enable Rules Editor, page 19-12.

b. In the Interfaces field, specify the interface or interfaces on which you want to enable the Botnet Traffic Filter. Normally, you want to enable the Internet-facing interface only. To select the interfaces or interface role objects using the Interfaces Selector, click **Select** (see Understanding Interface Role Objects, page 6-73).

You can configure a global classification that applies to all interfaces by selecting the All Interfaces role object (selected by default). If you configure an interface-specific classification, the settings for that interface overrides the global setting.

c. Do one of the following to identify the traffic that you want to monitor:

- To monitor all traffic, leave the ACL field blank.
- To specify the traffic that you want to monitor, click **Select** to the right of the ACL field to select an Access Control List object that identifies the traffic that you want to monitor. For example, you might want to monitor all port 80 traffic on the outside interface. For more information about Access Control List objects, see Creating Access Control List Objects, page 6-53.

Note You can specify only one enable rule per interface.

d. Click **OK**.

The BTF Enable Rules Editor closes and the rule is added to the Enable Rules table.

Step 3 To automatically drop malware traffic, follow these steps:

Note You must enable the Botnet Traffic Filter for the traffic you want to automatically drop before creating a drop rule for that traffic.

a. On the Traffic Classification tab, click **Add Row** under the Drop Rules table.

This opens the BTF Drop Rules Editor, page 19-13.

b. In the Interfaces field, specify the interface or interfaces on which you want to drop traffic. There must be a corresponding enable rule for the interface. To select the interfaces or interface role objects using the Interfaces Selector, click **Select** (see Understanding Interface Role Objects, page 6-73).

You can configure a global classification that applies to all interfaces by selecting the All Interfaces role object (selected by default). If you configure an interface-specific classification, the settings for that interface overrides the global setting.

c. Do one of the following to identify the traffic that you want to drop:

- To monitor all traffic, leave the ACL field blank.

- To specify the traffic that you want to monitor, click **Select** to the right of the ACL field to select an Access Control List object that identifies the traffic that you want to monitor. For example, you might want to monitor all port 80 traffic on the outside interface. For more information about Access Control List objects, see Creating Access Control List Objects, page 6-53.

d. In the Threat Level area, choose one of the following options to drop traffic specific threat levels. The default level is a range between Moderate and Very High.

Note We highly recommend using the default setting unless you have strong reasons for changing the setting.

- Value—Specify the threat level you want to drop.
- Range—Specify a range of threat levels.

Note Static block list entries are always designated with a Very High threat level.

e. Click **OK**.

The BTF Drop Rules Editor closes and the rule is added to the Drop Rules table.

Step 4 To add more rules, repeat steps 2 and 3, as required. When finished adding rules, click **Save** to save your changes.

Step 5 To treat graylisted traffic as blocked traffic for action purposes, on the Dynamic Blacklist Configuration tab, check the **Treat Ambiguous traffic as Blacklist** check box.

If you do not enable this option, graylisted traffic will not be dropped if you configure a drop rule for that traffic.

Botnet Traffic Filter Rules Page

You can use the Botnet Traffic Filter Rules page to define rules for identifying malicious traffic passing through your ASA security device.

The Botnet Traffic Filter Rules page is divided into three sections:

- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11
- Whitelist/Blacklist Tab, page 19-14

Navigation Path

To access the Botnet Traffic Filter Rules page, do one of the following:

- (Device view) Select a device, then select **Firewall > Botnet Traffic Filter Rules** from the Policy selector.
- (Policy view) Select **Firewall > Botnet Traffic Filter Rules** from the Policy Type selector. Select an existing policy or create a new one.
- (Map view) Right-click a device and select **Edit Firewall Policies > Botnet Traffic Filter Rules**.

Related Topics

- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11
- BTF Enable Rules Editor, page 19-12
- BTF Drop Rules Editor, page 19-13
- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Configure DNS Dialog Box, page 17-19

Dynamic Blacklist Configuration Tab

Use the Dynamic Blacklist Configuration tab to enable database updates from the Cisco update server and to enable use of the downloaded dynamic database by the security appliance.

Navigation Path

From the Botnet Traffic Filter Rules Page, page 19-9, click the **Dynamic Blacklist Configuration** tab.

Related Topics

- Configuring the Dynamic Database, page 19-4
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9
- Traffic Classification Tab, page 19-11
- BTF Enable Rules Editor, page 19-12
- BTF Drop Rules Editor, page 19-13
- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Configure DNS Dialog Box, page 17-19

Field Reference

Table 19-1 Dynamic Blacklist Configuration Tab

Element	Description
Enable Dynamic Blacklist From Server	Enables downloading of the dynamic database from the Cisco update server. If you do not have a database already installed on the security appliance, it downloads the database after approximately 2 minutes. The update server determines how often the security appliance polls the server for future updates, typically every hour. **Note** If the device is in multiple context mode, configure this option on the System context for that device.

Table 19-1 Dynamic Blacklist Configuration Tab (continued)

Element	Description
Use Dynamic Blacklist	Enables use of the dynamic database for the Botnet Traffic Filter. **Note** In multiple context mode, you configure use of the database on a per-context basis.
Treat Ambiguous traffic as Blacklist	When selected, graylisted traffic will be treated as blocked traffic for action purposes. If you do not enable this option, graylisted traffic will not be dropped if you configure a drop rule for that traffic.

Traffic Classification Tab

Use the Traffic Classification tab to view or to configure the traffic classification definitions for a device or shared policy and to identify malicious traffic that you want automatically dropped. Traffic classification definitions (enable rules) consist of an interface or interface role with an associated ACL that identifies the traffic that is monitored by the Botnet Traffic Filter. You can configure settings for specific interfaces or for interface roles. You can use the All Interfaces role object to enable botnet filtering globally (selected by default). If you configure an interface-specific classification, the settings for that interface override any settings defined for an interface role.

For a particular interface, you can specify only one enable rule that identifies the traffic that is subject to Botnet Traffic Filtering; however, you can specify multiple drop rules to identify traffic that should be dropped by the Botnet Traffic Filter.

Note We highly recommend configuring Dynamic Filter Snooping for proper functioning of the Botnet Traffic Filter. When in Device view, Cisco Security Manager provides a link at the bottom of the Traffic Classification tab that will take you directly to the Inspection Rules page so that you can enable Dynamic Filter Snooping. For more information, see Enabling DNS Snooping, page 19-6.

The columns in the tables summarize the settings for an entry and are explained in BTF Enable Rules Editor, page 19-12 and BTF Drop Rules Editor, page 19-13.

To configure traffic classification and actions:

- Click the **Add Row** button to add an interface or interface role to the table, and fill in the BTF Enable Rules Editor, page 19-12 or BTF Drop Rules Editor, page 19-13.
- Select an entry and click the **Edit Row** button to edit an existing entry.
- Select an entry and click the **Delete Row** button to delete it.

Navigation Path

From the Botnet Traffic Filter Rules Page, page 19-9, click the **Traffic Classification** tab.

Related Topics

- BTF Enable Rules Editor, page 19-12
- BTF Drop Rules Editor, page 19-13
- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Understanding Botnet Traffic Filtering, page 19-1

- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9
- Dynamic Blacklist Configuration Tab, page 19-10
- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Configure DNS Dialog Box, page 17-19

BTF Enable Rules Editor

Use the BTF Enable Rules Editor to specify the interfaces on which you want to enable the Botnet Traffic Filter and to identify the traffic that you want to monitor. You can specify only one enable rule per interface.

Navigation Path

To access the BTF Enable Rules Editor, right-click inside the work area of the Enable Rules table on the Traffic Classification tab and then select **Add Row**, or right-click an existing entry and select **Edit Row**.

Related Topics

- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9
- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11
- BTF Drop Rules Editor, page 19-13
- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Configure DNS Dialog Box, page 17-19

Field Reference

Table 19-2 BTF Enable Rules Editor

Element	Description
Interfaces	The interfaces or interface roles on which you want to enable the Botnet Traffic Filter. Enter the name of the interface or the interface role, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. You can use the All Interfaces role object to enable botnet filtering globally (selected by default). If you configure an interface-specific classification, the settings for that interface override the global settings. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. See Understanding Interface Role Objects, page 6-73.

Table 19-2 BTF Enable Rules Editor (continued)

Element	Description
ACL	Specifies the access-list to use for identifying the traffic that you want to monitor. If you do not specify an access list, by default you monitor all traffic. To specify the traffic that you want to monitor, click **Select** to the right of the ACL field to select an Access Control List object that identifies the traffic that you want to monitor. For example, you might want to monitor all port 80 traffic on the outside interface. For more information about Access Control List objects, see Creating Access Control List Objects, page 6-53.

BTF Drop Rules Editor

Use the BTF Drop Rules Editor to identify malware traffic that you want to automatically drop. You can specify multiple drop rules per interface.

Navigation Path

To access the BTF Drop Rules Editor, right-click inside the work area of the Drop Rules table on the Traffic Classification tab and then select **Add Row**, or right-click an existing entry and select **Edit Row**.

Related Topics

- Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9
- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11
- BTF Enable Rules Editor, page 19-12
- Whitelist/Blacklist Tab, page 19-14
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Configure DNS Dialog Box, page 17-19

Field Reference

Table 19-3 BTF Drop Rules Editor

Element	Description
Interfaces	The interfaces or interface roles on which you want to enable the Botnet Traffic Filter. Enter the name of the interface or the interface role, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. You can use the All Interfaces role object to enable botnet filtering globally (selected by default). If you configure an interface-specific classification, the settings for that interface override the global settings. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. See Understanding Interface Role Objects, page 6-73.
ACL	Specifies the access-list to use for identifying the traffic that you want to monitor. If you do not specify an access list, by default you monitor all traffic. To specify the traffic that you want to monitor, click **Select** to the right of the ACL field to select an Access Control List object that identifies the traffic that you want to monitor. For example, you might want to monitor all port 80 traffic on the outside interface. For more information about Access Control List objects, see Creating Access Control List Objects, page 6-53.
Threat Level	The Threat Level fields identify the threat level of malicious traffic that you want dropped. The default level is a range between Moderate and Very High. **Note** We highly recommend using the default setting unless you have strong reasons for changing the setting. • Value—Specify the threat level you want to drop. – Very-low – Low – Moderate – High – Very-high • Range—Specify a range of threat levels. **Note** Static block list entries are always designated with a Very High threat level.

Whitelist/Blacklist Tab

Use the Whitelist/Blacklist tab to view or to configure the static database entries for a device or shared policy. The Device Blacklist contains domain names or IP addresses of malicious or undesirable sites. You can use the static block list to supplement the Cisco dynamic database or you can use the static block list alone if you can identify all the malware sites that you want to target.

The Device Whitelist contains domain names or IP addresses of sites that are deemed to be acceptable. If the dynamic database includes blocked addresses that you think should not be blocked, you can manually enter them into a static allow list. Static allow list entries take precedence over entries in the static block list and the Cisco dynamic database. Addresses on the allow list still generate syslog messages, but because you are only targeting blocked syslog messages, they are informational.

To configure the static database:

- Click the **Add Row** button to define static database entries using the Device Whitelist or Device Blacklist Dialog Box, page 19-15.
- Select an entry and click the **Edit Row** button to edit an existing entry.

Timesaver Select an entry and press **F2** or double-click on an entry in the Device Whitelist or Device Blacklist to edit that entry in place.

- Select an entry and click the **Delete Row** button to delete it.

Navigation Path

From the Botnet Traffic Filter Rules Page, page 19-9, click the **Whitelist/Blacklist** tab.

Related Topics

- Adding Entries to the Static Database, page 19-5
- Understanding Botnet Traffic Filtering, page 19-1
- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Device Whitelist or Device Blacklist Dialog Box, page 19-15
- Botnet Traffic Filter Rules Page, page 19-9
- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11

Device Whitelist or Device Blacklist Dialog Box

Use the Device Whitelist or Device Blacklist dialog box to manually define domain names or IP addresses that you want to add to the allow lists (safe) or block lists (malicious). You can use the static block list to supplement the Cisco dynamic database or you can use the static block list alone if you can identify all the malware sites that you want to target. Names or addresses that appear on both the allow list and the dynamic block list are identified only as allowed addresses in syslog messages and reports.

Domain names can be complete (including the host name, such as www.cisco.com), or partial (such as cisco.com). For partial names, all web site hosts on that domain are either allowed or blocked. You can also enter host IP addresses. Use a comma or new line to separate multiple entries.

Navigation Path

From the Whitelist/Blacklist Tab, page 19-14, click the **Add Rows** button beneath the Device Whitelist or Device Blacklist tables, or select an entry and click the **Edit Row** button.

Related Topics

- Adding Entries to the Static Database, page 19-5
- Understanding Botnet Traffic Filtering, page 19-1

- Task Flow for Configuring the Botnet Traffic Filter, page 19-2
- Botnet Traffic Filter Rules Page, page 19-9
- Dynamic Blacklist Configuration Tab, page 19-10
- Traffic Classification Tab, page 19-11
- BTF Enable Rules Editor, page 19-12
- BTF Drop Rules Editor, page 19-13
- Whitelist/Blacklist Tab, page 19-14
- Configure DNS Dialog Box, page 17-19

Working with ScanSafe Web Security

Security Manager provides integration with ScanSafe Web Security. ScanSafe Web Security is a cloud-based SaaS (Security as a Service) function that makes available its web security data centers at various locations worldwide. When ScanSafe Web Security is integrated with a router, selected HTTP and HTTPS traffic is redirected to ScanSafe Cloud for content scanning and for malware detection by other means. Also, you can use ScanSafe Web Security to provide differentiated services to particular users, user groups, and IPs.

Invoking ScanSafe Web Security from Security Manager, you can define policies and settings in the following areas:

- Content scanning settings
- Content scanning policies
- AAA server settings
- AAA policies

With ScanSafe Web Security integration in Security Manager you can copy and share most policies and framework-based policy features. The following table details the scope of support for scanning and AAA policy types.

Supported Type	Example
Content Scan Settings	Primary server IP, secondary server IP, server timeouts
Content Scan Policies	Global whitelist policies Include/exclude user groups, default user configuration, default user group configuration Interfaces that must have content scanning enabled
AAA Server Settings	Identity policy object used in http-basic and NTLM policy Http-basic and NTLM related timeouts Order of occurrence for proxy, http-basic, and NTLM LDAP server and LDAP attribute map configuration for IOS **Note** Radius and TACAS servers are also supported. Per interface AAA list
AAA Policies	Http-basic and NTLM admission rule support (authentication methods) now join the previously available Auth-Proxy method

Security Manager does *not* support the following features:

- PAM configuration when inspect or ZBF rules for http/https are not present
- Auth-proxy using LDAP on older IOS versions. (That is, only IOS versions that support ScanSafe Web Security)
- Identity policy with auth-proxy as AAA method. (Support only for NTLM and http-basic.)
- Validation of Virtual Template number for identity policy creation
- Validation of the Secure Trust Point for LDAP server
- Inheritance of content scanning rules
- AD browsing of user groups and users
- Tool support for newer policies (such as policy query)
- Control tag policy

For more information on the ScanSafe Web Security product, go to http://www.cisco.com/en/US/partner/products/ps11720/index.html.

Related Topics

- ScanSafe Web Security Page, page 20-4
- ScanSafe Web Security Settings Page, page 20-6
- Configuring ScanSafe Web Security, page 20-2
- Add and Edit Default User Groups Dialog Box, page 20-6

This chapter contains the following topics:

- Configuring ScanSafe Web Security, page 20-2
- ScanSafe Web Security Page, page 20-4
- ScanSafe Web Security Settings Page, page 20-6

Configuring ScanSafe Web Security

Use the ScanSafe Web Security Settings page to define the settings for the default user group. As with other settings policies, you can share the default user group policy settings.

Related Topics

- ScanSafe Web Security Page, page 20-4
- ScanSafe Web Security Settings Page, page 20-6
- Chapter 20, “Working with ScanSafe Web Security”
- Add and Edit Default User Groups Dialog Box, page 20-6
- AAA Rules Page, page 15-10

Note All steps are shown as performed from the Policy view.

To configure ScanSafe Web Security, perform the following steps:

Step 1 From the Policy Types selector, select Firewall > ScanSafe Web Security.

The ScanSafe Web Security page appears with the Interfaces tab selected.

Step 2 Enable those interfaces by which web requests are to be forwarded to the ScanSafe Web Security server by selecting them from the list in the **Available Interfaces** column and moving them to the **Selected Interfaces** column.

Step 3 Select the **WhiteListing Regular Expressions** tab.

Step 4 Select the **Notify Tower** checkbox to send notifications to the ScanSafe Web Security server regarding the whitelisting. It is applicable to all whitelisting except that which is IP-based.

(ScanSafe Web Security receives a warning when no regular expression is specified for white listing.)

Step 5 In the HTTP Host area specify the regular expressions to be whitelisted (using regular expression matching) by selecting them from the list in the **Available Regular Expressions** column and moving them to the **Selected Regular Expressions** column.

Step 6 In the HTTP User Agent area specify the regular expressions to be whitelisted by selecting them from the list in the **Available Regular Expressions** column and moving them to the **Selected Regular Expressions** column.

Step 7 Select the **WhiteListing ACLs** tab.

Step 8 Specify the type of ACLs to operate upon by selecting either **Extended** or **Standard** from the Type list.

Step 9 Specify the ACLs to whitelist by selecting them from the list in the column on the left and moving them to the **Selected items** column.

Step 10 Select the **User Groups** tab.

Tip You can use the User Groups page to define user groups, specify both the default user and default user group, and to include or exclude user groups. You can also edit or delete entries in all three of these lists.

Step 11 Specify a default user by entering the user name in the **Default User** field (optional).

Step 12 Specify a default user group by entering the user group name in the **Default User Group** field.

Step 13 Include a user group by selecting the interface and then adding the user group to the Include list.

Step 14 Exclude a user group by selecting the interface and then adding the user group to the Exclude list.

Step 15 Select **Policy > Firewall > Settings > ScanSafe Web Security** from the policy selector.

Step 16 With the **Details** tab selected, specify the Primary ScanSafe Server by entering the following values:

- IP Address/Name
- HTTP Port (default 8080)
- HTTPS Port (default 8080)

Step 17 With the **Details** tab selected, specify the Secondary ScanSafe Server by entering the following values:

- IP Address/Name (only a valid IP address or FQDN).
- HTTP Port (default 8080)
- HTTPS Port (default 8080)

Step 18 Specify the **Server Timeout** period in seconds (default 300).

Step 19 Specify the **Session Idle Timeout** period in seconds (default 300).

Step 20 Specify the source address by doing *one* of the following:

- Click the **IP Address** button and then enter the IP address.

- Click the **Interface** button, and then click the **Select** button and browse the Interface Selector to select an interface.

Note A valid source IP or interface must be one of the interfaces on which ScanSafe Web Security is enabled (on the Firewall > ScanSafe Web Security page > Interface tab.

Step 21 Enter the **License** and select the checkbox if it is encrypted.

Tip If Encrypted is not selected, the value entered must be 32 hexadecimal characters.

Step 22 If desired, select the **Enable Logging** checkbox.

ScanSafe Web Security Page

Security Manager provides integration with ScanSafe Web Security. ScanSafe Web Security is a cloud-based SaaS (Security as a Service) function that makes available its web security data centers at various locations worldwide. When ScanSafe Web Security is integrated with a router, selected HTTP and HTTPS traffic is redirected to ScanSafe Cloud for content scanning and for malware detection by other means. Also, you can use ScanSafe Web Security to provide differentiated services to particular users, user groups, and IPs.

Using ScanSafe Web Security in Security Manager, you can define settings and policies in the following areas:

- Content scanning settings
- Content scanning policies
- AAA server settings
- AAA policies

With ScanSafe Web Security integration in Security Manager you can copy and share most policies and framework-based policy features.

For information on how best to configure the ScanSafe Web Security Page for your particular purposes, see Working with ScanSafe Web Security.

Navigation Path

(Policy view) Select Firewall and open Settings from the Policy Type selector. Then click ScanSafe Web Security to open the ScanSafe Web Security Settings Page.

Note Configuration of the ScanSafe Web Security policies and settings is also possible by way of the Map view.

Related Topics

- Chapter 20, “Working with ScanSafe Web Security”
- Configuring ScanSafe Web Security, page 20-2
- ScanSafe Web Security Settings Page, page 20-6

- Add and Edit Default User Groups Dialog Box, page 20-6
- AAA Rules Page, page 15-10

Field Reference

Element	Description
Interfaces Tab	
—Filter	Details on using filters in Security Manager are found at Filtering Tables, page 1-48.
Interfaces	This tab allows you to select interfaces and Security Manager-defined interface roles on which web requests will be forwarded to the ScanSafe Web Security server for content scanning.
—Available Interfaces	Interfaces that are available to be selected for ScanSafe Web Security.
—Selected Interfaces	Interfaces selected must be facing the WAN on which hosts' requests for web services are forwarded to ScanSafe Web Security server
-	-
Whitelisting Regular Expressions Tab	
—Notify Tower	This checkbox, when selected, specifies that the ScanSafe Web Security tower must be notified regarding the whitelisting. It is applicable to all ACL-based whitelisting variants except IP based whitelisting. The default behavior is that the notification is not sent.
—Available Regular Expressions (HTTP Host)	Lists the regular expressions available and considered for delivery to the ScanSafe Web Security server.
——Filter (HTTP Host)	Enables the administrator to filter whitelisted regular expressions sent to ScanSafe Web Security server by specifying include and exclude user group list. It operates on a match-all or match-any basis.
—Selected Regular Expressions (HTTP Host)	A host that matches the selected regular expressions is whitelisted, and is not redirected to the ScanSafe Web Security server.
—Available Regular Expressions (HTTP User Agent)	An agent that matches the available regular expressions is whitelisted, and is not redirected to the ScanSafe Web Security server.
—Selected Regular Expressions (HTTP User Agent)	When configured, only regular expressions that are in the Selected Regular Expressions list are sent to ScanSafe Cloud.
Whitelisting ACLs Tab	
—ACL Type	Specifies the type of ACL Whitelisting, either standard or extended. **Note** Standard ACLs used for whitelisting are discovered as extended ACLs. A prefix of "CSM_EXT_" is added to the ACL name. Standard ACls are converted to extended ACLs as extended ACLs are complete and recommended
—Selected ACLS	When configured, only regular expressions that are in the Selected Regular Expressions list are sent to ScanSafe Cloud.

Element	Description
User Groups Tab	
—Default User	A global name that is sent to the ScanSafe Web Security server when there is no content-scan-session specific user name. Use it when you want the same content scan policy for all users in a branch office (for example).
—Default User Group	A global name that is sent to the ScanSafe Web Security server when there is no content-scan-session specific user name. Use it when you want the same content scan policy for all user groups in a branch office (for example).
—Interface Specific Default User Groups	Lists default user group for each interface.
—Include/Exclude	You can use the Include and Exclude lists to specify the particular user groups to be included or excluded.

Add and Edit Default User Groups Dialog Box

Use the Default User Groups dialog box to specify the default user group for a particular interface.

This dialog box is only useful when the ISR cannot determine user credentials, but would like to assign the users (that is, the IP addresses) to a user group so that other group-based policies in the ISR can be enforced. Only one group can be configured on an interface.

For details on these ScanSafe Web Security server configuration settings, see the ScanSafe Web Security Settings Page.

Related Topics

- ScanSafe Web Security Page, page 20-4
- ScanSafe Web Security Settings Page, page 20-6
- Chapter 20, “Working with ScanSafe Web Security”
- Configuring ScanSafe Web Security, page 20-2
- AAA Rules Page, page 15-10

Navigation Path

(Policy view) Select Firewall and open the ScanSafe Web Security Page. Then click on the User Groups tab.

ScanSafe Web Security Settings Page

Related Topics

- ScanSafe Web Security Page, page 20-4
- Chapter 20, “Working with ScanSafe Web Security”
- Configuring ScanSafe Web Security, page 20-2
- Add and Edit Default User Groups Dialog Box, page 20-6

- AAA Rules Page, page 15-10

Navigation Path

(Policy view) Select Firewall and open Settings from the Policy Type selector. Then click ScanSafe Web Security to open the ScanSafe Web Security Settings Page.

(Device view) Select Firewall and open Settings from the Policy Type selector. Then click ScanSafe Web Security to open the ScanSafe Web Security Settings Page.

Field Reference

Table 20-1 ScanSafe Web Security Settings

Element	Description	Usage
IP Address|Name (Primary ScanSafe Server)	The primary FQDN or IP address of the server configured to operate ScanSafe Web Security.	Both
HTTP Port (Primary ScanSafe Server)	Default primary port for proxied HTTP traffic (default=8080).	Both
HTTPS Port (Primary ScanSafe Server)	Default primary port for proxied HTTPS traffic (default=8080).	Both
IP Address/Name (Backup ScanSafe Server)	The secondary FQDN or IP address of the server configured to operate ScanSafe Web Security.	Both
HTTP Port (Backup ScanSafe Server)	Default secondary port for proxied HTTP traffic (default=8080).	Both
HTTPS Port (Secondary ScanSafe Server)	Default secondary port for proxied HTTPS traffic (default=8080).	Both
Server Timeout	Polling timeout when checking the availability of the ScanSafe Web Security server.	IOS Only
Session Idle Timeout	Inactivity timeout of the ScanSafe Web Security server (default=300 seconds). Used to remove the session if it is found inactive.	IOS Only
On Failure	Determines the action to be taken (Drop all Traffic or Allow All Traffic) when both primary and secondary ScanSafe Web Security servers are found inactive.	IOS Only
IP Address (Source Address)	IP address from which a packet to the ScanSafe Web Security server originates from the router.	IOS Only
Interface (Source Address)	Interface address from which a packet to the ScanSafe Web Security server originates from the router.	IOS Only
License	The license sent to the ScanSafe Web Security server (32 hexadecimal characters).	Both
Encrypted	When selected, enables the encryption. ASA does not accept encrypted license text to be configured.	IOS Only
Enable Logging Checkbox	Enables IOS syslogs (default=not enabled).	IOS Only
Public Key File	Name of the public key file	ASA Only

Table 20-1 ScanSafe Web Security Settings (continued)

Element	Description	Usage
Connection Retry Count	Number of times that the system should retry connecting.	ASA Only

Managing Zone-based Firewall Rules

The Zone-based Firewall feature (also known as Zone-based Policy Firewall) allows unidirectional application of IOS firewall policies between groups of interfaces known as "zones." That is, interfaces are assigned to zones, and firewall rules are applied to specific types of traffic moving in one direction between the zones. Beginning with 4.16, Cisco Security Manager allows sub-interfaces to be assigned to zones. Zone-based firewalls enforce a secure inter-zone policy by default, meaning traffic cannot pass between security zones until an explicit policy allowing that traffic is defined.

The "zone" itself is an abstraction—multiple interfaces with the same or similar security requirements that can be logically grouped together. For example, router interfaces Ethernet 0/0 and Ethernet 0/1 might be connected to the local LAN. When viewed from a firewall perspective, these two interfaces are similar in that they represent the internal network, and they can be grouped into a single zone for the purposes of firewall configuration. Then you can specify firewall policies between that and other zones. These inter-zone policies offer considerable flexibility and granularity, so different inspection policies can be applied to multiple host groups connected to the same router interface.

Note The zone-based firewall feature is supported on IOS devices running 12.4(6)T or later, and ASR devices running 12.2(33) or later.

A Simple Example

A security zone should be configured for each region of similar security within the network, so that all interfaces assigned to the same zone are protected with a similar level of security. For example, consider an access router with three interfaces:

- One interface is connected to the public Internet
- One interface is connected to a private LAN that must not be accessible from the public Internet
- One interface is connected to an Internet-service "demilitarized zone" (DMZ), where a Web server, Domain Name System (DNS) server, and e-mail server must have access to the public Internet

Each interface in this network would be assigned to its own zone, as shown in the following figure.

Figure 21-1 Basic Security Zone Topology

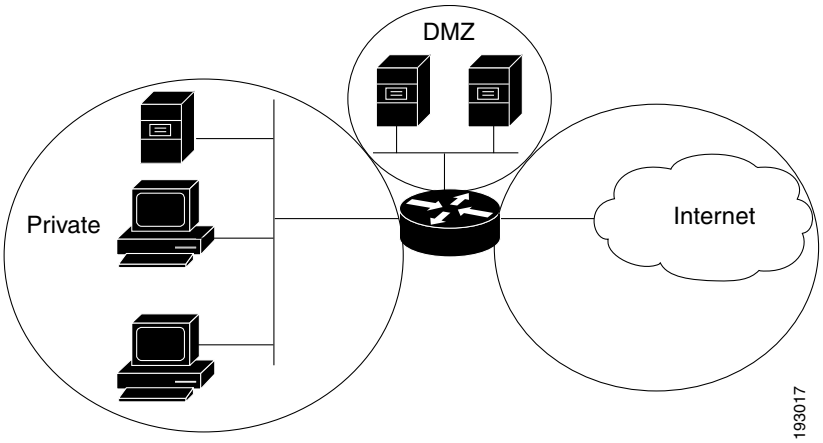

This example configuration typically would have three main policies (sets of rules) defining:

- Private zone connectivity to the Internet
- Private zone connectivity to DMZ hosts
- Internet zone connectivity to DMZ hosts

Zone-based firewalls impose a prohibitive default security posture. In other words, for example, unless the DMZ hosts are specifically allowed access to other networks, those networks are protected against any undesired connections from the DMZ hosts. Similarly, unless access is specifically provided for Internet hosts to access the Private zone directly, the Private zone hosts are safe from unwanted access by Internet hosts.

In this simple example, each zone has only one member interface. If an additional interface is added to the Private zone, for example, the hosts connected to that new interface can immediately pass traffic to all hosts connected to the existing interface in the zone. Additionally, traffic to hosts in other zones is immediately controlled by existing Private zone policies.

In a more realistic example, you might allow varied access from the public Internet to specific hosts in the DMZ, and varied application-use policies for hosts in the protected LAN.

This chapter contains the following topics:

- Understanding the Zone-based Firewall Rules, page 21-3
- Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules, page 21-8
- Understanding the Relationship Between Services and Protocols in Zone-based Firewall Rules, page 21-11
- General Recommendations for Zone-based Firewall Rules, page 21-12
- Developing and Applying Zone-based Firewall Rules, page 21-12
- Adding Zone-Based Firewall Rules, page 21-13
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Changing the Default Drop Behavior, page 21-48
- Configuring Settings for Zone-based Firewall Rules, page 21-49
- Troubleshooting Zone-based Rules and Configurations, page 21-54

- Zone-based Firewall Rules Page, page 21-58

Understanding the Zone-based Firewall Rules

Zones establish the security borders of your network. A zone defines a boundary where traffic is subjected to inspection or filtering as it crosses to another region of your network. The default zone-based firewall policy between zones is "deny all." Thus, if no zone-based firewall rules are explicitly configured, all traffic moving between all zones is blocked.

Zone-based firewall rules apply specific actions—Drop, Pass, Inspect, and Content Filter—to various types of unidirectional traffic between pairs of zones. The direction of the traffic is determined by specifying a source and destination zone as part of each rule.

Logging

Zone-based firewall rules offer syslog, alert, and audit-trail logging options. Most messages are logged to the router console unless a syslog server is configured. See Logging on Cisco IOS Routers, page 65-1 for information about configuring syslog logging.

Important Points

Please note the following points regarding zones and zone-based firewall rules:

- Zone-based firewall rules are supported only on IOS devices running 12.4(6)T or later, and ASR devices running 12.2(33) or later.
- If a zone-based firewall rule and an IOS Inspection rule use the same interface, an error results.

  The zone-based firewall model and the earlier interface-based inspection rules model are not mutually exclusive on the router, but they cannot be combined on any given interface. That is, an interface cannot be configured as a member of a security zone if it is configured with Inspection rules. Further, configuring a router to use both models at the same time is not recommended.
- An interface can be assigned to only one security zone, but zones can include multiple interfaces. If an interface is assigned to more than one zone, an error results.
- All traffic to and from a given interface is implicitly blocked when the interface is assigned to a zone (except traffic to and from other interfaces in the same zone, and traffic to any interface on the router). Thus, to permit traffic to and from a zone-member interface, one or more rules allowing or inspecting traffic must be configured between that zone and any other zone.
- Traffic is implicitly allowed to flow between interfaces that are members of the same zone. However, you can define rules that require inspection of traffic between same-zone members.
- The "Self" zone is a default zone that defines the router itself as a separate security zone, which you can specify as either the source or destination zone. The Self zone is the only exception to the default "deny all" policy. All traffic to any router interface is allowed until explicitly denied.

  A zone-based firewall rule that includes the Self zone applies to local traffic—that is, traffic directed to the router, or to traffic generated by the router; it does not apply to traffic through the router. See The Self Zone, page 21-5 for more information.
- The Inspect action is not allowed in rules that apply to the Self zone.
- The Pass action permits traffic in one direction only. You must explicitly define rules for return traffic. However, with the Inspect action, return traffic is automatically allowed for established connections.
- Traffic cannot flow between a zone-member interface and any interface that is not a zone member.

- Interfaces that have not been assigned to a zone can still function as classical router ports and might still have other types of firewall rules configured on them.

  However, if an interface is not part of your zone-based firewall policy, it might still be necessary to add that interface to a zone and configure a “pass all” policy (sort of a “dummy policy”) between that zone and any other zone to which inter-zone traffic flow is desired.
- Access-control list (ACL) rules applied on interfaces that are also zone members are processed before the zone rules are applied. Therefore, to continue using both rule types, it may be necessary to relax the interface ACLs to ensure certain traffic flows are processed by the zone-based rules.
- All interfaces in a zone must belong to the same Virtual Routing and Forwarding (VRF) instance. Zone-based rules can be configured between zones whose member interfaces are in separate VRFs. However, if traffic cannot flow between these VRFs, these rules will never be executed. See Zones and VRF-aware Firewalls, page 21-7 for more information.
- Zones are defined using Interface Role objects. If you change the definition of an interface role that you are using for a zone, you are changing the zone, which can affect existing traffic flows. In addition, if you use wildcards in the interface role to specify an interface name pattern, be aware that interfaces may automatically be added to the zone when you create new interfaces on the router.
- If zone-based firewall rules contain conflicting zone information, the first rule defined in the table takes precedence. Rules that do not reference valid zones are not deployed and an activity validation warning is shown.
- Empty zones result in activity validation errors for certain devices; refer to the following restriction lists.
- Source and destination zones cannot be the same for certain devices; refer to the following restriction lists.
- For devices that support nested class-maps (IOS devices running 12.4(6)T+ and ASR devices running 3.5S+/15.2(1)S+), Security Manager will configure multiple traffic classes as a single traffic class when appropriate.
- When discovering policies on an ISR, Security Manager will discover the zone-based firewall policies as they are on the device. When discovering policies on an ASR, Security Manager will try to match policies to existing object-groups and reuse them when possible.

ASR Restrictions

The following are restrictions specific to ASR devices:

- Deep Packet Inspection (DPI) is not allowed.
- Source and destination zones can be the same. This is possible because intra-zone traffic inspection is allowed.
- Content (URL) Filtering is not allowed.
- Only certain protocols are supported, such as DNS, FTP, H.323, ICMP, RTSP, SIP, Skinny, TCP, TFTP, and UDP.
- Although IOS-XE only supports security group tags, Security Manager allows the use of security group names for ASR devices, either directly or as part of a security group object group, and will automatically convert security group names to their tag number during deployment. However, if the ISE is not reachable, Security Manager will not be able to resolve the security group names. In such situations, Security Manager will not generate the rule for the configuration and an appropriate warning/error message will be displayed.

- IOS-XE only supports the use of security group object groups. Security Manager will generate a validation error message if inline security group tags or names are deployed to an ASR. Also, when discovering policies on an ASR, Security Manager will try to match security group object groups to existing object groups and reuse them when possible.

ISR Restrictions

The following are restrictions specific to ISR devices:

- Empty zones cannot exist.
- Source and destination zones cannot be the same.
- Cisco TrustSec Support for IOS feature is supported on the Cisco Integrated Services Router Generation 2 (ISR G2) only.
- Destination security group tags are not supported.
- Although IOS only supports security group tags, Security Manager allows the use of security group names for ISR devices and will automatically convert security group names to their tag number during deployment. However, if the ISE is not reachable, Security Manager will not be able to resolve the security group names. In such situations, Security Manager will not generate the rule for the configuration and an appropriate warning/error message will be displayed.
- Although IOS does not support security group object groups, Security Manager allows the use of security group object groups for ISR devices and will automatically expand the object groups to their tag numbers during deployment.

Related Topics

- The Self Zone, page 21-5
- Using VPNs with Zone-based Firewall Policies, page 21-6
- Zones and VRF-aware Firewalls, page 21-7
- Configuring Settings for Zone-based Firewall Rules, page 21-49
- Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules, page 21-8
- Understanding the Relationship Between Services and Protocols in Zone-based Firewall Rules, page 21-11
- General Recommendations for Zone-based Firewall Rules, page 21-12
- Developing and Applying Zone-based Firewall Rules, page 21-12

The Self Zone

The router itself is defined as a separate security zone, with the fixed name **Self**, and since IOS firewalls support examination of traffic (TCP, UDP and H.323 only) that terminates or originates on the router (together known as “local” traffic), incoming and outgoing router traffic can be subject to rules in the same way as routed inter-zone traffic.

When an interface is assigned to a zone, the hosts connected to that interface are included in that zone. By default, traffic is allowed to flow between interfaces that are members of the same zone, while a default “deny-all” policy is applied to traffic moving between zones.

However, traffic flowing directly between other zones and the router’s IP interfaces (the Self zone) is implicitly allowed. This ensures that connectivity to the router’s management interfaces is maintained when a zone firewall configuration is applied to the router.

This also means that traffic flowing to and from the IP addresses of the router's interfaces is not initially controlled by zone policies. If you wish to control traffic moving between the router interfaces and other zones, you must apply rules that block or allow this local traffic.

When configuring the rules for the Self zone, consider the following:

- All IP addresses configured on the router belong to the Self zone, regardless of interface zone membership.
- Traffic to and from the Self zone is unrestricted until you configure explicit rules to the contrary.

  That is, when you configure a zone-based firewall rule that includes the Self zone, traffic between the Self zone and the other zone is immediately restricted in both directions. For example, if you define a rule affecting traffic from the "Private" zone to the Self zone, the router cannot originate any traffic to the Private zone until you define one or more rules for Self to Private.

  Traffic between the router itself and other zones that are not included in the Self-zone rules remains unaffected.
- The Inspect action is not allowed in rules that apply to the Self zone.

When configuring restrictions on inbound Self-zone traffic, consider the necessary outbound traffic (including the routing and network management protocols). For example, if you restrict inbound traffic from a zone to the router itself, the routing protocols could stop working on all interfaces belonging to that zone.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3

Using VPNs with Zone-based Firewall Policies

Recent enhancements to the IP Security (IPsec) VPN implementation simplify firewall policy configuration for VPN connectivity. IPSec Virtual Tunnel Interface (VTI) and GRE+IPSec allow the confinement of VPN site-to-site and client connections to a specific security zone by placing the tunnel interfaces in that security zone. Connections can be isolated in a VPN DMZ if connectivity must be limited by a specific policy. Or, if VPN connectivity is implicitly trusted, VPN connections can be placed in the same security zone as the trusted inside network.

To configure the router to use zone-based firewall rules with dynamic VPNs (those which dynamically create Tunnel/Loopback/Virtual interfaces):

- Define a zone specifically for the VPN interfaces.
- Enter this zone in the **VPN Zone** field on the VPN tab of the Zone Based Firewall Page, page 21-50.
- Create zone-based firewall rules to allow the VPN traffic, as appropriate.

If non-VTI IPsec is employed, you must exercise caution when you configure a zone-based firewall policy for VPN. The zone policy must specifically allow access to protected hosts by remote VPN hosts or clients if they are in a different zone than the ingress interface for encrypted VPN traffic. This access policy must be configured by including an access control list (ACL) enumerating the source IP addresses of the VPN clients, and the destination IP addresses of all protected hosts the VPN clients are allowed to reach. If the access policy is not properly configured, the policy could expose vulnerable hosts to hostile traffic.

Refer to this white paper on cisco.com "Using VPN with Zone-Based Policy Firewall" for further discussion of these topics.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3

Zones and VRF-aware Firewalls

Cisco IOS firewalls are VRF-aware (Virtual Routing and Forwarding), providing management of IP address overlap across different VRFs, separate thresholds and timeouts for VRFs, and so forth. For application of zone-based firewall rules, all interfaces in a zone must belong to the same VRF.

When multiple VRFs are configured on a router and one interface provides common services to all the VRFs (for example, Internet service), you should place that interface in a separate zone. You can then define policies between the common zone and other zones. (There can be one or more zones per VRF.)

You can configure rules between two zones that contain different VRFs, as shown in the following illustration.

Figure 21-2 Zones and VRF

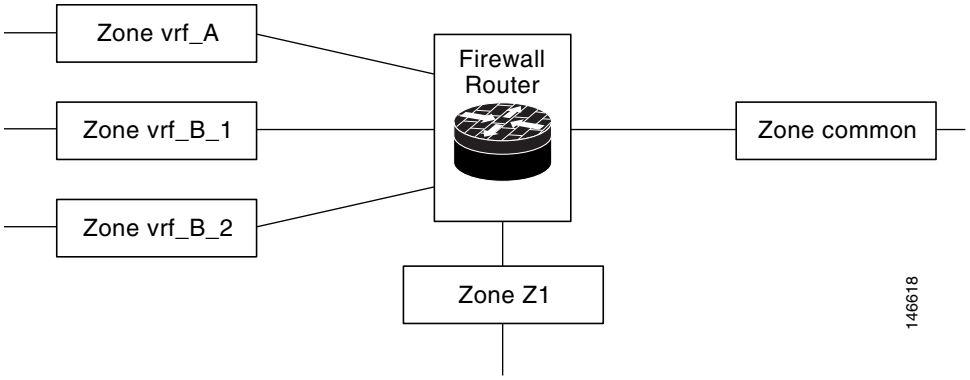

In this illustration:

- The interface providing common services is a member of the zone “common.”
- All of VRF A is in a single zone, “vrf_A.”
- VRF B, which has multiple interfaces, is partitioned into two zones “vrf_B_1” and “vrf_B_2.”
- Zone Z1 does not have VRF interfaces.

Based on this configuration:

- You can specify policies between each of these zones and the common zone. Additionally, you can specify polices between each of the zones vrf_A, vrf_B_n and Z1 if VRF route export is configured and the traffic patterns make sense.
- You can configure a policy between zones vrf_A and vrf_B_1, but be sure that traffic can flow between them.
- You do not need to specify the global thresholds and timers on a per-VRF basis. Instead, parameters are supplied to the Inspect action through a parameter map.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3

Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules

When you create a zone-based firewall rule, you must specify two execution-related settings: Permit/Deny and an Action (Drop, Pass, Inspect, or Content Filter). To obtain the results you want, you must clearly understand the relationship between these two parameters:

- **Permit/Deny**—The Permit/Deny setting appears to correspond to Permit/Deny in an access control list (ACL) entry. However, in zone-based firewall rules, unlike in standard access rules, these keywords do not permit or deny traffic. Instead, they specify whether you want to apply an Action to the traffic flow defined by the Source, Destination, and Services fields, and they affect processing of related class maps.
  - **Permit** – Applies the specified Action to traffic that matches the Source, Destination, and Services fields. (If protocols are listed in the Protocols table, the Action is further limited to those protocols.)

    Tip: Essentially all of your zone-based rules should be Permit rules. This is the easiest configuration to understand—it means you are identifying the traffic to which you want the chosen Action applied.
  - **Deny** – Exempts the traffic defined by the Source, Destination, and Services fields. (If protocols are listed in the Protocols table, the exemption is further limited to those protocols.) In other words, treat the traffic as not matching the rule. Instead, evaluate subsequent class maps (which are not the same as zone rules) for the zone pair and look for a subsequent map that matches the traffic. If no subsequent map matches the traffic, apply the default rule to the traffic (see Changing the Default Drop Behavior, page 21-48).

    It is important to understand that there is not a one-to-one relationship between zone rules and class maps. Therefore, you cannot determine just by looking at the rules table how the rules will be converted to class maps. You must preview the configuration to see which subsequent rules might be applied to traffic that matches your Deny rule. (To preview the configuration, save your changes and select **Tools > Preview Configuration**. For more information, see Previewing Configurations, page 8-44.)

    In general, you might use a Deny rule to exempt a specific IP address within a subnet from a Permit rule you want to apply to the subnet in general; for example, exempting 10.100.10.1 from a rule applying to 10.100.10.0/24. However, it is much easier to create a Permit rule for the specific IP address and apply a desired Action, and ensure that the rule is listed above the general rule in the zone-based rules table.

    If you decide to use Deny rules, be sure to also read Troubleshooting Zone-based Rules and Configurations, page 21-54.
- **Action**—The Action parameters define what happens to traffic that matches a Permit rule. These parameters are ignored for Deny rules, except to determine to which class map the rule is added.

  When you create a Permit rule, traffic that matches the Source, Destination, Services, and Protocol fields is processed according to the Action you choose: drop the traffic (and optionally log it), pass the traffic (and optionally log it), inspect the traffic, or apply content filtering (for Web traffic only).

  When you inspect traffic for some protocols, or perform content filtering, you have the option of specifying a policy map to use for deep inspection. The deep inspection policy map also specifies actions based on the deeper characteristics of the traffic. This additional inspection applies to packets that meet the requirements of the class map to which the assigned policy map refers. Packets that do not match the deep inspection class map are allowed. Thus, deep inspection might reset TCP connections if the policy map specifies that action.

The following table illustrates the relationship between Permit/Deny and the chosen Action in a zone-based firewall rule. The table uses the TCP service as an example, but the general explanation applies to the IP service as well. The result applies only to the From and To zones specified in the rule.

Table 21-1 Relationship Between Permit/Deny and Action in Zone-based Rules

Permit / Deny	Service	Rule Action	Protocol	Result
Permit	TCP	Pass	(None)	Pass all TCP traffic.
Deny	TCP	Pass	(None)	Skip the rule and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Pass action is ignored.
Permit	TCP	Drop	(None)	Drop all TCP traffic.
Deny	TCP	Drop	(None)	Skip the rule and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Drop action is ignored.
Permit	TCP	Pass	DNS	Only DNS traffic passes. Other TCP traffic is handled by subsequent rules.
Permit	TCP	Drop	DNS	DNS traffic is dropped. Other TCP traffic is handled by subsequent rules.
Deny	TCP	Pass	DNS	Skip the rule for DNS traffic and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Pass action is ignored.
Deny	TCP	Drop	DNS	Skip the rule for DNS traffic and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Drop action is ignored.
Permit	TCP	Inspect	HTTP	Allow and inspect HTTP traffic. If you specify a policy map for deep inspection, the action from the policy map is applied to any packets that match deep inspection parameters (for example, reset the connection for protocol violations).

Table 21-1 Relationship Between Permit/Deny and Action in Zone-based Rules (continued)

Permit / Deny	Service	Rule Action	Protocol	Result
Deny	TCP	Inspect	HTTP	Skip the rule for HTTP traffic and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Inspect action is ignored. **Tip** If subsequent rules, or the class default, pass the traffic without inspection, you need to create a Permit/Pass rule in the other direction (or an access rule) to allow return traffic for the HTTP connection. If your intention is to prohibit HTTP connections, create a Permit/Drop rule instead of a Deny/Inspect rule.
Permit	TCP	Content Filter	HTTP	Allow and inspect HTTP traffic, and apply URL filtering maps to selectively permit or deny Web connections based on the Web sites requested. If you specify a policy map for deep inspection, the action from the policy map is applied to any packets that match deep inspection parameters (for example, reset the connection for protocol violations). Thus, traffic can be dropped either because the Web site is on the block list, or because the HTTP packets violate your deep inspection rules.

Table 21-1 Relationship Between Permit/Deny and Action in Zone-based Rules (continued)

Permit / Deny	Service	Rule Action	Protocol	Result
Deny	TCP	Content Filter	HTTP	Skip the rule for HTTP traffic and evaluate the next class map. Either a subsequent class map with a Permit rule is applied, or the class default rule is applied. The Content Filter action is ignored. **Tip** This type of rule can exempt the specified source/destination from content filtering if no subsequent class maps drop the traffic or apply content filtering. However, if you want to allow HTTP connections for this traffic, you must create a Permit/Inspect rule for the traffic.

Understanding the Relationship Between Services and Protocols in Zone-based Firewall Rules

When you create a zone-based firewall rule, there are two seemingly similar parameters which help identify the characteristics of the target traffic: Services and Protocols. The entries in these fields can provide very similar information, but it is used differently when constructing zone-based firewall policies in the device configuration. This section describes the recommended uses of these fields.

- **Services** – The Services field is used to define traffic protocol(s) in an access control list (ACL) entry. Along with the Sources and Destinations specified, this ACL entry is used by a class map to define the traffic to which you want to apply a policy. However, unlike in standard access rules, the Services information is not the primary means of identifying the traffic protocol. It is required only because ACLs must have a service designation for each entry.

  In general, you can leave the default entry (IP) in the Services field for all of your zone-based firewall rules, using the Protocol table to identify specific protocols that you want to Drop, Pass, or Inspect.

  If you do elect to specify a Service other than IP, ensure that your selection does not conflict with any protocols listed in the Protocol table. For example, do not specify UDP in the Services field, and then list a TCP-based protocol in the table. In general, for a given rule, if you specify a specific service in the Services field, do not enter any protocols in the Protocol table.

- **Protocol** – The Protocol table, in the Action area of the Add and Edit Zone Based Rule dialog boxes, is used to select one or more protocols, add custom port application mappings (if you specify non-default ports), and apply deep inspection policy maps. You can specify very specific protocols, such as DNS, general protocols such as TCP and UDP, and even custom protocols that identify ports you use for special applications.

As a general rule, leave Services set to IP and use the Protocol table to identify the protocols (which are also services) for all of your zone-based rules for the Drop, Pass, and Inspect actions. (The Content Filter action automatically uses the HTTP protocol, which you can configure but not change.) Following this approach will create a configuration that is as "clean" and easy to interpret (and troubleshoot) as possible.

For more detailed information on how these fields are used when generating device configurations, see Troubleshooting Zone-based Rules and Configurations, page 21-54.

General Recommendations for Zone-based Firewall Rules

Zone-based firewall rules allow a wide variety of configurations. You can quickly generate a set of rules that will be very complex and difficult to analyze, because you can use the zone-based rules in place of the standard access rules, inspection rules, and Web filter rules.

When defining zone-based rules, strive to keep them as simple and straightforward as possible. Consider the following recommendations for helping to maintain simplicity in your zone-based firewall policy:

- Only use **Permit** rules. The chosen Action determines what happens to matched traffic, and **Deny** rules are difficult to analyze. For more information, see Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules, page 21-8.
- The **Drop** and **Pass** rules are equivalent to standard interface access rules, but are applied to the specified zone pair. You can use either the Services field or the Protocol table to identify the type of traffic, but we recommend using the Protocol table exclusively. To drop traffic, specify **Permit** with the Action **Drop**.
- You do not need to first pass traffic before inspecting it. For example, if you want to allow HTTP traffic between zones, you need only a single Permit/Inspect rule; you do not need to first create a Permit/Pass rule. If you do use Pass rules, note that you must also create a Pass rule in the return direction if you want to allow returning traffic. In practice, you generally can avoid creating Pass rules, using only Inspect rules.
- You can use Permit/Pass and Permit/Drop rules to perform the same functions as standard access rules. Thus, you can eliminate your access rules policy and use only zone-based firewall rules.

  However, because there are several tools available for analyzing interface access rules, and Security Manager allows you to use the same interface roles in zone-base rules and access rules, you might find it more convenient to create your Pass/Drop policies (which are Permit/Deny in standard access rules) in the access rules table instead of the zone rules table. Use the zone rules table primarily for zone-based Inspection and Content Filter rules.
- Use sections to organize the rules for each zone pair. Sections make it easy for you to see all of the rules for a pair, which can be critical if your rules have sequential dependencies. For more information on working with sections, see Using Sections to Organize Rules Tables, page 12-20.

Developing and Applying Zone-based Firewall Rules

The following is a general overview of how to develop and apply zone-based firewall rules to your network.

- Consider your network, and its sub-networks, in terms of security zones—think about the security requirements of the various zones. As a general guideline, group router interfaces that are similar when viewed from a security perspective.

- Determine the types of traffic to be examined as it travels from one zone to another, decide how each type is to be examined and handled.
- Define zone-based firewall rules that implement these decisions. This process may include some or all of the following procedures, which you can perform prior to defining the rules themselves, or which you can perform as necessary during rule definition:
  - Define the zones by creating named Interface Role objects, assigning the appropriate interfaces and interface patterns to them.
  - Define/edit Port Application Mapping (PAM) settings for specific Layer 4 protocols and ports, and optionally specific networks and hosts.
  - Configure Deep Packet Inspection (DPI) policies for Layer 7 protocols—HTTP, IMAP, instant messaging (IM), and peer-to-peer (P2P).
  - Configure Protocol Info parameter maps; these define DNS servers that interact with the IM applications.
  - Configure Inspect parameter maps that define connection, timeout, and other settings for the Inspect action.
  - Define WebFilter parameter or policy maps for URL-based content filtering.

The following topics provide additional information about these procedures:

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16

Adding Zone-Based Firewall Rules

This procedure explains how to configure a zone-based firewall rule in Security Manager.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Configuring Settings for Zone-based Firewall Rules, page 21-49
- Understanding Map Objects, page 6-78
- Enabling and Disabling Rules, page 12-20
- Adding and Removing Rules, page 12-9
- Moving Rules and the Importance of Rule Order, page 12-19

Step 1 Access the Zone-based Firewall Rules Page, page 21-58 as follows:

- (Device view) Select an IOS router and then select **Firewall > Zone Based Firewall Rules** from the Policy selector.
- (Policy view) Select **Firewall > Zone Based Firewall Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Click the Add Row button below the rules table, or right-click anywhere inside the table, and choose **Add Row** to open the Add Zone Based Firewall Rule dialog box.

Refer to Adding and Editing Zone-based Firewall Rules, page 21-62 for a complete description of this dialog box.

Step 3 Define the base Traffic flow for this rule.

Note Together, the Permit/Deny, Sources, Destinations, and Services options can be thought of as defining a simple access rule that can be enhanced by the application of in-depth Action-related policies, and restricted to a specific direction between a specific pair of zones.

a. Choose whether to Permit or Deny further processing of traffic that matches this rule. See Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules, page 21-8 for more information.

b. Optionally, provide Source and Destination hosts/networks or Security Groups (IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+ only).

   By default, the traffic definition encompasses packets from "any" source, to "any" destination. You can use these fields to refine this base traffic definition by providing one or more source and destination hosts/networks. (See Understanding Networks/Hosts Objects, page 6-80 and Selecting Security Groups in Policies, page 14-16 for more information.)

c. Specify one or more Services (protocols) that indicate the type of traffic; for example, IP, TCP, etc.

   You can provide more than one Service; however, IP generally stands alone. (See Understanding and Specifying Services and Service and Port List Objects, page 6-100.)

d. Provide the From Zone; that is, the only zone from which matched traffic can originate.

e. Provide the To Zone; that is, the only zone to which matched traffic can flow.

   Refer to Understanding Interface Role Objects, page 6-73 for more information about zone/interface objects.

   Note Together, the From Zone and the To Zone constitute what is sometimes referred to as a "zone-pair."

f. Click the Advanced button to add a time range, or to apply a packet-fragment or an established-connection restriction to this zone-based firewall rule.

   See Zone-based Firewall Rule: Advanced Options Dialog Box, page 21-67 for more information about these options.

Step 4 Specify the actions to be applied to traffic matching this definition by choosing a base Action, and supplying additional parameters as necessary.

a. Choose a base Action:

   - **Drop** – Matching traffic is silently dropped; no notification of the drop is sent to the originating host.
   - **Drop and Log** – Matching traffic is dropped and a syslog message generated; no notification of the drop is sent to the originating host.
   - **Pass** – Traffic is forwarded. This action is unidirectional; Pass allows traffic in only the specified direction.
   - **Pass and Log** – Traffic is forwarded and a syslog message generated.

Note The Pass actions do not track the state of connections or sessions within the traffic. Pass only allows the traffic in one direction. A corresponding rule must be defined to allow return traffic. The Pass actions are useful for protocols such as IPSec ESP, IPSec AH, ISAKMP, and other inherently secure protocols with predictable behavior. However, most application traffic is better handled in the zone-based firewall rules with the Inspect action.

- **Inspect** – This option offers state-based traffic control—the device maintains connection or session information for TCP and UDP traffic, meaning return traffic in reply to connection requests is permitted.

  Choose this option to apply packet inspection based on your selected Layer 4 (TCP, UDP) and Layer 7 (HTTP, IMAP, instant messaging, and peer-to-peer) protocols. You also can edit the Port Application Mapping (PAM) settings for the selected protocols, and you can set up deep packet inspection (DPI) and provide additional protocol-related information for the Layer 7 protocols.

- **Content Filter** – Lets you configure HTTP content inspection (URL filtering) based on a WebFilter parameter map, or a WebFilter policy map. This action is generally equivalent to a Web Filter rule; however, zone-based firewall rules support additional advanced options, such as HTTP deep packet inspection (DPI).

  The router intercepts HTTP requests, performs protocol-related inspection, and optionally contacts a third-party server to determine whether the requests should be allowed or blocked. You can provide a WebFilter parameter map, which defines filtering based on local URL lists, as well as information from an external SmartFilter (previously N2H2) or Websense server. Alternately, you can provide a WebFilter policy map that accesses Local, N2H2, Websense, or Trend Micro filtering data.

b. For any Action except Content Filter, you can select and edit the specific traffic Protocol(s) to be considered:

Click Select next to the Protocol table to open the Protocol Selector Dialog Box, page 21-68. Select one or more protocols and click >> to move them to the Selected Protocol list. You can edit the Port Application Mapping (PAM) settings for the selected protocols; see Configure Protocol Dialog Box, page 21-69 for more information.

The Instant Messaging and Stun-ice protocols also allow selection of Protocol Info parameter maps. Further, when Inspect is the chosen Action, some protocols allow selection of deep-inspection policy maps.

See Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16, and Configuring Protocol Info Parameter Maps, page 21-33 for more information.

Note It is not necessary to specify protocols for the Drop, Drop and Log, Pass, and Pass and Log actions. You can leave the Protocol table empty and pass or drop traffic based on the Sources, Destinations, and Services parameters.

c. When the chosen Action is Content Filter, configure the URL filtering:

1. Click Configure next to the Protocol field to customize the HTTP PAM settings, and to apply an HTTP deep-inspection policy map. See the Configure Protocol Dialog Box, page 21-69 for more information
2. Select either WebFilter Parameter Map, or WebFilter Policy Map, and enter or Select the name of the appropriate WebFilter map. See Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36 for more information.

d. When the chosen Action is Inspect or Content Filter, you can enter or Select the name of an Inspect Parameter map to apply a customized set of connection, timeout, and other settings. See Configuring Inspect Parameter Maps, page 21-31 for more information.

Step 5 (Optional) Enter a description to help you identify the rule.

Step 6 (Optional) Under Category, select a category to help you identify this rule in the rules table. See Using Category Objects, page 6-13.

Step 7 Click **OK** to close the Add Zone Based Firewall Rule dialog box and return to the Zone Based Firewall Rules table.

The new rule is listed in the table.

Configuring Inspection Maps for Zone-based Firewall Policies

When you configure zone-based firewall policies for a router, you can define rules to inspect traffic by choosing Inspect as the Action for the rule. You can then select the specific protocols to inspect.

For some protocols, you can select policy maps to perform deep inspection on packets that match your criteria. You can configure these maps from the policy object selector dialog box while defining the rule, or at any time in the Policy Object Manager window (select **Manage > Policy Objects**). In addition to policy maps, there are some parameter maps you can configure for inspection.

- For protocols that allow deep inspection, you can select a related policy map, which in turn incorporates class maps that define match conditions for the targeted traffic. To create these policy maps in the Policy Object Manager, select one of the available map types (which are listed in the following table) from the **Maps > Policy Maps > Inspect** folder, and review the detailed usage information in Configuring Policy Maps for Zone-Based Firewall Policies, page 21-34.

  For information on creating class maps for use in your deep-inspection policy maps, see the references to the match criterion dialog boxes in the following table, as well as the topic Configuring Class Maps for Zone-Based Firewall Policies, page 21-19. These class map are found in the **Maps > Class Maps > Inspect** folder in the Policy Object Manager.

- When Inspect (or Content Filter) is the chosen Action, you can also apply an Inspect Parameters map in the Adding and Editing Zone-based Firewall Rules. Zone-based firewall inspection includes several general settings, all of which have default values that are appropriate for most networks. If you want to adjust any of these settings, you must create an Inspect Parameters map. In the Policy Object Manager, select **Maps > Parameter Maps > Inspect > Inspect Parameters** and review the detailed usage information in Configuring Inspect Parameter Maps, page 21-31.

Table 21-2 Policy Objects for Zone-based Firewall Inspection Rules

Protocol	Minimum IOS Software Version	Policy Map	Class Map	Parameter Map	Description and Match Criteria Reference
Instant Messaging: AOL, ICQ, MSN Messenger, Windows Messenger, Yahoo Messenger	12.4(9)T	IM (Zone based IOS)	AOL ICQ MSN Messenger Windows Messenger Yahoo Messenger	Protocol Info	Inspect traffic based on the type of service (text-chat or any other). See Zone-based Firewall IM Application Class Maps: Add or Edit Match Condition Dialog Boxes, page 21-21. You must also select a Protocol Info parameter map to define the DNS servers used by the traffic you are inspecting. See Configuring Protocol Info Parameter Maps, page 21-33.
Peer-to-peer (P2P): eDonkey, FastTrack, Gnutella, Kazaa2	12.4(9)T	P2P	eDonkey FastTrack Gnutella Kazaa2	None	Inspect traffic based on file name. See Zone-based Firewall P2P Application Class Maps: Add or Edit Match Condition Dialog Boxes, page 21-21.
H.323	12.4(6)T	H.323 (IOS)	H.323 (IOS)	None	Inspect traffic based on the H.323 message type. See H.323 (IOS) Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-22.
HTTP	12.4(6)T	HTTP (Zone based IOS)	HTTP (IOS)	None	Inspect traffic based on a wide variety of criteria including the content of the header or body, port misuse, and whether the traffic includes a Java applet. See HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes, page 21-22.
IMAP (Internet Message Access Protocol) POP3 (Post Office Protocol 3)	12.4(6)T	IMAP POP3	IMAP POP3	None	Inspect traffic based on invalid commands or clear-text logins. See IMAP and POP3 Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-25.

Table 21-2 Policy Objects for Zone-based Firewall Inspection Rules (continued)

Protocol	Minimum IOS Software Version	Policy Map	Class Map	Parameter Map	Description and Match Criteria Reference
SIP (Session Initiation Protocol)	12.4(6)T	SIP (IOS)	SIP (IOS)	None	Inspect traffic based on a wide variety of criteria. See SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes, page 21-25.
SMTP (Simple Mail Transfer Protocol)	12.4(6)T	SMTP	SMTP	None	Inspect traffic based on data length. See SMTP Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-27.
Stun-ice	12.4(9)T	None	None	Protocol Info	You must select a Protocol Info parameter map to define the DNS servers used by the traffic you are inspecting. See Configuring Protocol Info Parameter Maps, page 21-33.
Sun RPC (Remote Procedure Call)	12.4(6)T	Sun RPC	Sun RPC	None	Inspect traffic based on the RPC protocol number. See Sun RPC Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-29.
SCTP (Stream Control Transmission Protocol)	12.4(6)T *Verify this row.*	SCTP	None	None	Inspect traffic based on PPID match criterion. See SCTP Policy Maps Add or Edit Match Condition and Action Dialog Boxes, page 17-92
Diameter protocol	12.4(6)T *Verify this row.*	Diameter	Diameter	None	Inspect traffic based on application ID, command codes, and AVP. See Diameter Class and Policy Maps Add or Edit Match Condition (and Action) Dialog Boxes, page 17-95
LISP (Locator and ID Separation Protocol)	12.4(6)T *Verify this row.*	LISP	None	None	Inspect traffic based on application ID, command codes, and AVP.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Zone-based Firewall Rules Page, page 21-58

- Creating Policy Objects, page 6-9
- Understanding Map Objects, page 6-78

Configuring Class Maps for Zone-Based Firewall Policies

Use the Add and Edit Class Map dialog boxes to define class maps to be used in policy maps of the same type. The name of the dialog box indicates the type of map you are creating.

A class map defines application traffic based on criteria specific to the application. You then select the class map in the corresponding policy map and configure the action to take for the selected traffic. Thus, each class map must contain traffic that you want to handle in the same way (for example, to allow it or to drop it).

When configuring zone-based firewall rules for devices running Cisco IOS Software, you can create class maps for the following purposes:

- For 12.4(6)T and higher, you can create classes for the inspection of the following types of traffic: H.323, HTTP, IMAP, POP3, SIP, SMTP, and Sun RPC. You can create classes for web filtering using the following class types: Local, N2H2 (SmartFilter), and WebSense. See the following topics for information on the match criteria:
  - H.323 (IOS) Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-22
  - HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes, page 21-22
  - IMAP and POP3 Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-25
  - SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes, page 21-25
  - SMTP Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-27
  - Sun RPC Class Maps Add or Edit Match Criterion Dialog Boxes, page 21-29
  - Local Web Filter Class Add or Edit Match Criterion Dialog Boxes, page 21-29
  - N2H2 and Websense Class Add or Edit Match Criterion Dialog Boxes, page 21-30
- For 12.4(9)T and higher, you can create classes for the inspection of the following types of traffic: AOL, eDonkey, FastTrack, Gnutella, ICQ, Kazaa2, MSN Messenger, Windows Messenger, and Yahoo Messenger. See the following topics for information on the match criteria:
  - Zone-based Firewall IM Application Class Maps: Add or Edit Match Condition Dialog Boxes, page 21-21
  - Zone-based Firewall P2P Application Class Maps: Add or Edit Match Condition Dialog Boxes, page 21-21
- For 12.4(20)T and higher, you can create classes for web filtering using the Trend policy object. Match criteria for Trend Content Filter class maps is described in the table below.

Navigation Path

Select **Manage > Policy Objects**, then select any zone-based class map object in the folders in the **Maps > Class Maps** folder in the table of contents. Right-click inside the work area, then select **New Object**, or right-click a row, then select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36

- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-3 Add or Edit Class Maps Dialog Boxes for Zone-Based Firewall Policies

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Match table Match Type (Except for Trend Content Filter class maps.)	The Match table lists the criteria included in the class map. Each row indicates whether the inspection is looking for traffic that matches or does not match each criterion and the criterion and value that is inspected. The name of the table indicates whether every one of the criteria must be met for the traffic to match the class (Match All), or whether matching any of the listed criteria is sufficient (Match Any). For the HTTP (IOS) and SMTP classes, you can choose whether to match all or any. When using a Match All table, if you add more than one criteria, ensure that you are not defining a set of characteristics that no traffic can match. **Tip** Match All works for devices running Cisco IOS Software version 12.4(20)T or higher only. • To add a criterion, click the Add button and fill in the Match Criterion dialog box. For more information, see the topics referenced above. • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.
Trend Content Filter Match Criteria	The match criteria for Trend Content Filter class maps differs from that of all other class maps. Instead of adding items to a table, you simply select the items you want from a list. Select the Enable checkbox for any of the Trend-Micro classifications on the following tabs. Traffic matches the class if it matches any of your selections. • Productivity Categories—Matches the traffic to the category to which the URL belongs. For example, you can target traffic associated with gambling or pornography. • Security Ratings—Matches the traffic to the security rating assigned to it by Trend-Micro. For example, you can target adware, which is traffic associated with advertising. See the Trend-Micro documentation for specific information on these categories or security classifications.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 21-3 Add or Edit Class Maps Dialog Boxes for Zone-Based Firewall Policies (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Zone-based Firewall IM Application Class Maps: Add or Edit Match Condition Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the various instant messenger (IM) application classes used with zone-based firewall policies to define a match criterion and value for the class map.

You can define a match for the following types of traffic:

- Any—Any type of traffic from the application except text chat traffic.
- Text-chat—Text chat traffic.

Navigation Path

From the Add or Edit Class Maps dialog boxes for AOL, ICQ, MSN Messenger, Windows Messenger, or Yahoo Messenger classes, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Zone-based Firewall P2P Application Class Maps: Add or Edit Match Condition Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the various peer-to-peer (P2P) application classes used with zone-based firewall policies to define a match criterion and value for the class map.

Navigation Path

From the Add or Edit Class Maps dialog boxes for eDonkey, FastTrack, Gnutella, or Kazaa2 classes, right-click inside the table and select **Add Row**, or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-4 Zone-based Firewall P2P Application Class Maps Add or Edit Match Condition Dialog Boxes

Element	Description
Criterion	Choose which criterion to match: • File Transfer – Matches file-transfer traffic. • Search Filename – Matches the names of files for which the user is searching. You can use this criterion to block users from searching for particular files using eDonkey. • Text Chat – Matches eDonkey text chat traffic.
Type	Specifies that the map includes traffic that matches the criterion.
File Name	The name of the file associated with the traffic. You can use regular expressions to specify a name pattern. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions, page 17-109. **Tip** eDonkey does not require a file name.

H.323 (IOS) Class Maps Add or Edit Match Criterion Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the Add or Edit Match Criterion dialog boxes for the H.323 (IOS) class used with zone-based firewall policies to define a match criterion and value for the class map. You can match traffic based on the H.323 protocol message type. Select the message that you want to match.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the H.323 (IOS) class, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the Add or Edit Match Criterion dialog boxes for the HTTP (IOS) class used with zone-based firewall policies to define a match criterion and value for the class map.

The fields on this dialog box change based on the criterion you select. You can use the following criteria:

- Request/Response Body Length, Request Body Length, Response Body Length—Specifies that the body length of the request, response, or both, is less than or greater than the specified number. This allows you to set a minimum or maximum message length.
- Request/Response Body, Request Body, Response Body—Applies a regular expression to match the body of the request, response, or both.
- Request/Response Header, Request Header, Response Header—You can match a regular expression against the header, test for repeated fields, check the content type, or check the total length or number of records in the header.
- Request/Response Protocol Violation—Matches non-compliant HTTP traffic.
- Request Argument, Request URI—Matches the length or content (with a regular expression) of the argument (parameters) or uniform resource identifier (URI) in a request message.
- Request Port Misuse—Matches the misuse of ports by certain types of applications.
- Response Body Java Applet—Matches Java applets in an HTTP connection.
- Response Header Status Line—Applies a regular expression to match the content of the status line in the header.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the HTTP (IOS) class, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-5 HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes

Element	Description
Criterion	Specifies which criterion of HTTP traffic to match. The criteria are described above.
Type	Specifies that the map includes traffic that matches the criterion.
Variable Fields	
The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Less Than Length	The minimum length in bytes of the evaluated field. The criterion matches if the length is less than the specified number.
Greater Than Length	The maximum length in bytes of the evaluated field. The criterion matches if the length is greater than the specified number.

Table 21-5 HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes (continued)

Element	Description
Header Option	The type of header record. If you do not select a record type, the count or expression is applied to all records in the header. If you select a record type, those selections are applied only to the records of the selected type. If you select content type or transfer encoding, you can make additional selections related to those types.
Request Method	The request method you want to match.
Value (Content Type)	If you select content-type in the Header Option field, you can select these types: • Mismatch—Verifies the content-type of the response message against the accept field value of the request message. • Unknown—The content type is not known. Select Unknown when you want to evaluate the item against all known MIME types. • Violation—The content-type definition and the content type of the actual body do not match.
Encoding Type	If you select transfer encoding in the Header Option field, you can select these types: • All—All of the transfer encoding types. • Chunked—The message body is transferred as a series of chunks; each chunk contains its own size indicator. • Compress—The message body is transferred using UNIX file compression. • Deflate—The message body is transferred using zlib format (RFC 1950) and deflate compression (RFC 1951). • GZIP—The message body is transferred using GNU zip (RFC 1952). • Identity—No transfer encoding is performed.
Greater Than Count	The maximum number of records allowed in the header. If you select a specific header option, the count applies to those types of records. If you do not select a specific header option, the count applies to the total number of records in the header without regard to type.
Regular Expression	The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object.
Port Misuse	The type of request port misuse you want to match. Your options are: • Any—Any of the listed types of misuse. • IM—Instant messaging protocol applications subject to inspection. • P2P—Peer-to-peer protocol applications subject to inspection. • Tunneling—Tunneling applications subject to inspection: HTTPPort/HTTPHost.

IMAP and POP3 Class Maps Add or Edit Match Criterion Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the Internet Message Access Protocol (IMAP) and Post Office Protocol 3 (POP3) classes used with zone-based firewall policies to define a match criterion and value for the class map.

You can select the following criteria to identify matching traffic:

- Invalid Command—Matches commands that are not valid on a POP3 server or IMAP connection.
- Login Clear Text—Matches non-secure logins, where the password is being provided in clear text.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the IMAP or POP3 classes, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use the Add or Edit Match Criterion dialog boxes for the SIP (IOS) class used with zone-based firewall policies to define a match criterion and value for the class map.

The fields on this dialog box change based on the criterion you select.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the SIP (IOS) class, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-6 SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes

Element	Description
Criterion	Specifies which criterion of traffic to match. You can select from the following: • Protocol Violation—Matches traffic that violates the protocol. • Request/Response Header Options—Matches a regular expression against the selected request or response header field. • Request Options—Matches the request method or matches a regular expression against the selected request header field. • Response Options—Matches a regular expression against the selected response header field or status message.
Type	Specifies that the map includes traffic that matches the criterion.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Header	The type of header in the request or response message. The regular expression is matched against the content of headers of the selected type.
Method	The request method you want to inspect: • ack—Acknowledges that the previous message is valid and accepted. • bye—Signifies the intention to terminate a call. • cancel—Terminates any pending request. • info—Communicates mid-session signaling information along the signaling path for the call. • invite—Sets up a call. • message—Sends an instant message. • notify—Informs subscribers of state changes. • options—Queries the capabilities of another user agent or a proxy server. • prack—Provides reliable transfer of provisional response messages. • refer—Indicates that the recipient should contact a third party using the contact information provided in the request. • register—Includes a contact address to which SIP requests for the address-of-record should be forwarded. • subscribe—Requests notification of an event or set of events at a later time. • update—Permits a client to update parameters of a session but has no impact on the state of a dialog.

Table 21-6 SIP (IOS) Class Add or Edit Match Criterion Dialog Boxes (continued)

Element	Description
Status	The regular expression is matched against the status line in the response.
Regular Expression	The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object.

SMTP Class Maps Add or Edit Match Criterion Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the SMTP classes used with zone-based firewall policies to define a match criterion and value for the class map.

Tip Only the Data Length criterion is available for routers running Cisco IOS Software lower than 12.4(20)T.

The fields on this dialog box change based on the criterion you select. You can use the following criteria:

- Data Length—Specifies that the data length of the traffic is greater than the specified number. You can match the data length of the traffic to determine if the data transferred in an SMTP connection exceeds the specified length in bytes. By default, inspection keeps data length below 20.
- Body Regular Expression—Applies a regular expression to match the content types and content encoding types for text and HTML in the body of an e-mail message. Only text or HTML that uses 7-bit or 8-bit encoding is checked. The regular expression cannot be scanned in messages that use another encoding type (such as base64 or zip files).
- Command Line Length—Specifies that the length of the ESMTP command line not be greater than the specified number. Use this to thwart Denial of Service (DoS) attacks.
- Command Verb—Limits inspection to the selected SMTP or ESMTP command. If you configure inspection for SMTP, all commands are inspected unless you limit them.
- Header Length—Specifies that the length of the SMTP header is greater than the specified number. Use this to thwart DoS attacks by limiting the possible size of the header.
- Header Regular Expression—Applies a regular expression to match the content of the header of an e-mail message. For example, you can use this to test for particular patterns in the subject, from, or to fields.
- Mime Content-Type Regular Expression—Applies a regular expression to match the Multipurpose Internet Message Exchange (MIME) content type of an e-mail attachment. Use this to prevent the transmission of undesired types of attachments.
- Mime Encoding—Specifies the MIME encoding type for e-mail attachments that you want to inspect. You can use this to identify unknown or non-standard encodings to restrict their transmission.
- Recipient Address—Applies a regular expression to match the recipient of an e-mail message in the SMTP RCPT command. Use this to search for a non-existent recipient, which might help you identify the source of spam.
- Recipient Count—Specifies that the number of recipients for an e-mail message cannot be greater than the specified number. Use this to prevent spammers from sending e-mails to a large number of users.

- Recipient Invalid Count—Specifies that the number of invalid recipients for an e-mail message cannot be greater than the specified number. Use this prevent spammers from sending e-mails to a large number common names, where they are fishing for real addresses. SMTP typically replies with a “no such address” message when an address is invalid; by putting a limit on the number of invalid addresses, you can prevent these replies to spammers.
- Reply EHLO—Specifies the service extension parameter in an EHLO server reply. Use this to prevent a client from using a particular service extension.
- Sender Address—Applies a regular expression to match the sender of an e-mail message. Use this to block specific senders, such as known spammers, from sending e-mail messages through the device.

Navigation Path

From the Add or Edit Class Maps dialog boxes for SMTP classes, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-7 SMTP Class Add or Edit Match Criterion Dialog Boxes

Element	Description
Criterion	Specifies which criterion of SMTP traffic to match. The criteria are described above.
Type	Specifies that the map includes traffic that matches the criterion.
Variable Fields The following fields vary based on what you select in the Criterion field. This list is a super-set of the fields you might see.	
Greater Than Length	The maximum length in bytes of the evaluated field. The criterion matches if the length is greater than the specified number.
Greater Than Count	The maximum number of recipients or invalid recipients allowed in the e-mail message. The criterion matches if the number is greater than the specified number.
Verb Option User Defined Format (For the Command Verb criterion.)	The SMTP or ESMTP command that you want to inspect. If you select User Defined, you must enter the text string that corresponds to a word in the body of the e-mail message. The word cannot include spaces or special characters; only alphanumeric characters.
Service Extension Parameter User Defined Format (For the Reply EHLO criterion.	The service extension parameter of an EHLO server reply that you want to inspect. Select one of the well-known parameters, or select User Defined to specify a private extension in the User Defined Format field.

Table 21-7 SMTP Class Add or Edit Match Criterion Dialog Boxes (continued)

Element	Description
Encoding Format User Defined Format	The MIME encoding format for which you want to test. Encoding types are: • 7-bit—ASCII encoding. • 8-bit—Used for the exchange of e-mail messages containing octets outside the 7-bit ASCII range. • base64—Encodes binary data by treating it numerically and translating it into a base 64 representation. • quoted-printable-Encoding that uses printable characters to transmit 8-bit data over a 7-bit data path. • binary—Encodes using only 0 and 1. • unknown—Encoding type is not known. • x-uuencode-Nonstandard encoding. • user defined—An encoding type you define. If you select User Defined, you must enter the text string that defines the encoding type you are looking for.
Regular Expression	The regular expression object that defines the regular expression you want to use for pattern matching. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new regular expression object.

Sun RPC Class Maps Add or Edit Match Criterion Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the Sun Remote Procedure Call (RPC) classes used with zone-based firewall policies to define a match criterion and value for the class map. You can enter the RPC protocol number that you want to match. See the Sun RPC documentation for information about protocol numbers.

Navigation Path

From the Add or Edit Class Maps dialog boxes for Sun RPC classes, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Local Web Filter Class Add or Edit Match Criterion Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the Local web filter class to define a match criterion and value for the class map.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the Local web filter class, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-8 Local Web Filter Class Add or Edit Match Criterion Dialog Boxes

Element	Description
Criterion	Specifies which criterion of traffic to match. You can select from the following: • Server Domain—Matches traffic based on the name of the server. The URLF Glob parameter map you select should specify server domain names such as *.cisco.com or www.cisco.com. • URL Keyword—Matches traffic based on keywords in the URLs. A key word is any complete string that occurs between / characters in a URL. For example, in the URL segment www.cisco.com/en/US, en and US are examples of keywords.
Type	Specifies that the map includes traffic that matches the criterion.
URLF Glob Parameter Map	The URLF Glob parameter map object that defines the URL patterns that you want to match. Ensure that the object you select has the appropriate content for the type of matching you selected. Enter the name of the object. You can click **Select** to choose the object from a list of existing ones or to create a new object.

N2H2 and Websense Class Add or Edit Match Criterion Dialog Boxes

Use the Add or Edit Match Criterion dialog boxes for the N2H2 (SmartFilter) and Websense web filter classes to define a match criterion and value for the class map. The only match criterion available is to match any response from the SmartFilter or Websense server.

Navigation Path

From the Add or Edit Class Maps dialog boxes for the N2H2 or Websense web filter class, right-click inside the table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Configuring Inspect Parameter Maps

Use the Add and Edit Inspect Parameter Map dialog boxes to define a parameter map for inspection for zone-based firewall policies on routers. If you configure the action of a zone-based firewall policy rule as Inspect or Content Filter, you can select an inspect parameter map to define connection, timeout, and other settings for the inspection action. If you do not select an inspect parameter map for a zone-based firewall rule, the system uses default values for these settings.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Inspect > Inspect Parameters** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-9 Add or Edit Inspect Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
DNS Timeout	The length of time, in seconds, for which a DNS lookup session is managed while there is no activity.
ICMP Timeout	The length of time, in seconds, for which an inactive ICMP (Internet Control Message Protocol) session is maintained.
Max Incomplete Low Max Incomplete High	The number of existing half-open sessions that will cause the software to start (at the high threshold) and stop (at the low threshold) deleting half-open sessions. Ensure that you enter a lower number in the Low field than you enter in the High field, for example, 400 and 500. The default is unlimited half-open sessions.
One Minute Low One Minute High	The number of new unestablished sessions that causes the system to start and stop deleting half-open sessions. Ensure that you enter a lower number in the Low field than you enter in the High field. The default is unlimited.
Max Sessions	The maximum number of inspection sessions on a zone pair, for example, 200. The default is unlimited.
TCP FINWAIT Timeout	How long to maintain TCP session state information after the firewall detects a FIN-exchange, in seconds. The FIN-exchange occurs when the TCP session is ready to close.

Table 21-9 Add or Edit Inspect Parameter Map Dialog Boxes (continued)

Element	Description
TCP SYNWAIT Timeout	How long to wait for a TCP session to reach the established state before dropping the session, in seconds.
TCP Idle Timeout	How long to maintain a TCP session while there is no activity in the session, in seconds.
TCP Max Incomplete Hosts TCP Max Incomplete Block Time	The threshold and blocking time (in minutes) for TCP host-specific denial-of-service (DoS) detection and prevention. The maximum incomplete hosts is the number of half-open TCP sessions with the same host destination address that can simultaneously exist before the software starts deleting half-open sessions to that host. An unusually high number of half-open sessions with the same destination host address could indicate that a DoS attack is being launched against the host. When the threshold is exceeded, half-open sessions are dropped based on the maximum incomplete block time: • If the block time is 0, the software deletes the oldest existing half-open session for the host for every new connection request to the host. This ensures that the number of half-open sessions to a given host never exceeds the threshold. • If the block time is greater than 0, the software deletes all existing half-open sessions for the host and then blocks all new connection requests to the host. The software continues to block all new connection requests until the block time expires. The software sends syslog messages whenever the specified threshold is exceeded and when blocking of connection initiations to a host starts or ends.
UDP Idle Timeout	How long to maintain a UDP session while there is no activity in the session, in seconds. When the software detects a valid UDP packet, the software establishes state information for a new UDP session. Because UDP is a connectionless service, there are no actual sessions, so the software approximates sessions by examining the information in the packet and determining if the packet is similar to other UDP packets (for example, it has similar source or destination addresses) and if the packet was detected soon after another similar UDP packet. If the software detects no UDP packets for the UDP session for the period of time defined by the UDP idle timeout, the software will not continue to manage state information for the session.
Enable Alert	Whether to generate stateful packet inspection alert messages on the console.
Enable Audit Trail	Whether audit trail messages are logged to the syslog server or router.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 21-9 Add or Edit Inspect Parameter Map Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring Protocol Info Parameter Maps

Use the Add and Edit Protocol Info Parameter Map dialog boxes to define a parameter map for the inspection of Instant Messaging (IM) applications or the Stun-ice protocol for zone-based firewall policies on routers. If you configure the action of a zone-based firewall policy rule as Inspect, you must select a protocol info parameter map when you configure any of these applications: AOL, ICQ, MSN Messenger, Windows Messenger, Yahoo Messenger, Stun-ice. The protocol info parameter map defines the DNS servers that interact with these applications, which helps the instant messenger application engine to recognize the instant messenger traffic and to enforce the configured policy for that instant messenger application.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Inspect > Protocol Info Parameters** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-10 Add or Edit Protocol Info Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
DNS Server Table	The DNS servers for which traffic will be permitted (and inspected) or denied. • To add servers, click the Add button and fill in the Add Server dialog box (see Add or Edit DNS Server for Protocol Info Parameters Dialog Box, page 21-34). • To edit a server, select it and click the Edit button. • To delete a server, select it and click the Delete button.

Table 21-10 Add or Edit Protocol Info Parameter Map Dialog Boxes (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit DNS Server for Protocol Info Parameters Dialog Box

Use the Add or Edit DNS Server dialog box to identify DNS servers for which traffic will be permitted (and inspected) or denied. These servers are defined in a Protocol Info parameter map for use with the inspection of protocols that require them in a zone-based firewall policy.

You can identify a server using any of these types:

- Server Name—The name of the DNS server. You can use an asterisk (*) as a wildcard character to match one or more characters. For example, if you want to identify all DNS servers on the cisco.com domain, you can specify *.cisco.com.
- IP Address—The IP address of a single DNS server.
- IP Address Range—A range of IP addresses identifying any DNS server within the start and end addresses.

Navigation Path

From the Add or Edit Protocol Info Parameter Map dialog boxes, click the **Add** button beneath the server table, or select a server and click the **Edit** button. See Configuring Protocol Info Parameter Maps, page 21-33.

Configuring Policy Maps for Zone-Based Firewall Policies

Use the Add and Edit Policy Map dialog boxes for zone-based firewall policies to define the match criterion and values for an inspection map used in a zone-based firewall policy for a Cisco IOS router. You can create policy inspection maps for H.323 (IOS), HTTP (Zone based IOS), IM (Zone based IOS), IMAP, P2P, POP3, SIP (IOS), SMTP, and Sun RPC inspection, and the name of the dialog box indicates the type of map you are creating.

When defining the inspection map, you select class maps of the same type and define the action to take for matching traffic. You can configure the required class maps before creating the policy maps or while you are creating them.

Navigation Path

Select **Manage > Policy Objects**, then any of the following items in the **Maps > Policy Maps > Inspect** folder in the table of contents: H.323 (IOS), HTTP (Zone based IOS), IM (Zone based IOS), IMAP, P2P, POP3, SIP (IOS), SMTP, and Sun RPC. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-11 Add or Edit Policy Maps Dialog Boxes for Zone-Based Firewall Policies

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Match All table	The Match All table lists class maps included in the policy map, and the action to take for traffic that matches the class. For traffic to match this class, all criteria defined in the selected class maps must be met. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see Add or Edit Match Condition and Action Dialog Boxes for Zone-Based Firewall and Web Filter Policies, page 21-35). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Match Condition and Action Dialog Boxes for Zone-Based Firewall and Web Filter Policies

Use the Add or Edit Match Condition and Action dialog boxes for zone-based firewall and web filter policies to select the class maps for inspection and to define the action to take for traffic that matches the class. This dialog box is used for the following types of policy maps: H.323 (IOS), HTTP (Zone based IOS), IM (Zone based IOS), IMAP, P2P, POP3, SIP (IOS), SMTP, Sun RPC, Web Filter.

The fields on this dialog box differ slightly depending on the type of policy map you are defining.

Navigation Path

From the Add or Edit Policy Maps dialog boxes for Zone-Based Firewall Policies, right-click inside the match table and select **Add Row** or right-click a row and select **Edit Row**. See Configuring Policy Maps for Zone-Based Firewall Policies, page 21-34.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-12 Add or Edit Match Condition and Action Dialog Boxes for Zone-Based Firewall Policies

Element	Description
Match Type	Indicates that you are selecting a class map. You must define class maps when creating policy maps for zone-based firewall policies.
Class Map P2P, IM, and Web Filter class map types.	The name of the class map for the type of policy map you are creating. Click **Select** to select the map from a list or to create a new class map object. For P2P, IM, and Web Filter policy maps, you must also select the type of policy map you are creating. For example, in a P2P map you must select between eDonkey, FastTrack, Gnutella, and Kazaa2. In an IM (Zone Based IOS) map, you must select between AOL, MSN Messenger, Yahoo Messenger, Windows Messenger, and ICQ. In a Web Filter map, you must select between Local, N2H2, WebSense, and Trend.
Action	The action you want the device to take for traffic that matches the selected class.

Configuring Content Filtering Maps for Zone-based Firewall Policies

When you configure zone-based firewall policies for a router, you can define rules to filter Web content by choosing Content Filter as the Action for the rule.

To filter Web content, you must configure certain map objects, which you can do from the policy object selector dialog box while defining the rule, or at any time in the Policy Object Manager window (select **Manage > Policy Objects**).

The type of maps required depends on the technique you are using to filter content, and on the Cisco IOS software version you are using. You can filter content based on URL lists defined locally on the device, or you can use external filtering servers such as SmartFilter (N2H2), Websense, or Trend Micro.

Tip If you use an external server, you must have set up and configured the server appropriately based on the documentation for the type of server you select. If you use Trend Micro servers, you must specify the server details, and register the product and download certificates, on the Content Filtering tab of the Zone Based Firewall page (select Firewall > Settings > Zone Based Firewall). See Zone Based Firewall Page, page 21-50.

The following are requirements for the map objects used with zone-based content filtering:

- For devices running releases below 12.4(20)T, you must create a URL Filter parameter map. In the Policy Object Manager, select **Maps > Parameter Maps > Web Filter > URL Filter**, and review the detailed usage information in Configuring URL Filter Parameter Maps, page 21-43.
  - To perform local filtering on the router using lists of allowed and denied hosts, create the lists on the Local Filtering tab. Any Web access request is first compared to these lists before the request is sent on to an external filtering server (if you have configured one). These lists contain either complete domain names (such as www.cisco.com), or partial names (such as cisco.com), but they do not include paths or page names, and you cannot use wildcards.
  - To use a SmartFilter (N2H2) or Websense server, configure the type of server you are using and its address information on the External Filter tab. You can also configure other settings that control communication with the server. You cannot configure a Trend Micro server using the URL Filter parameter map.
- For devices running release 12.4(20)T and higher, the preferred approach is to use a Web Filter policy map. Although Web Filter policy maps are more complex, they provide added flexibility, and they let you access Trend Micro filtering servers. In the Policy Object Manager, select **Maps > Policy Maps > Web Filter > Web Filter**, and review the detailed usage information in Configuring Web Filter Maps, page 21-47.

  A Web Filter policy map incorporates other types of maps. To create the policy map, you will need one or more of these other types of maps:

  - Parameter maps – On the Parameters tab of the Add and Edit Web Filter Map dialog boxes, you can select parameter maps for the various types of Web filtering if you do not want to use the default settings. If you are using SmartFilter (N2H2) or Websense, you need to select a parameter map because the map identifies those servers. For Local and Trend Micro filtering, parameter maps configure some general settings, the most interesting of which is whether to display a message or Web page when a URL is blocked. In the Policy Object Manager, you can find parameter maps for Local, N2H2, Trend, and Websense in the **Maps > Parameter Maps > Web Filter** folder. For detailed usage information, see Configuring Local Web Filter Parameter Maps, page 21-38, Configuring N2H2 or WebSense Parameter Maps, page 21-39, or Configuring Trend Parameter Maps, page 21-42.

    Note You configure Trend Micro server information on the Content Filtering tab of the Zone Based Firewall page (select Firewall > Settings > Zone Based Firewall). See Zone Based Firewall Page, page 21-50.

  - Class maps for match conditions – These class maps define the type of traffic you want to target and specify the action to be taken. You select a type of filtering (Local, SmartFilter/N2H2, Websense, or Trend Micro), specify the class map that identifies the targeted traffic, and choose an action (such as Allow, Reset, etc.) to be taken for that traffic. In the Policy Object Manager, you can find class maps for Local, N2H2, Trend, and Websense in the **Maps > Class Maps > Web Filter** folder.

    These class-map configurations depend on the type of filtering:

    Local Filtering – The Local WebFilter class map is a list of one or more URLF Glob parameter maps that specify either domain names or URL keywords that you want to target. A URL keyword is any text string delineated by forward-slash (/) characters in a URL. These class maps help you define allowed and denied URL lists for a WebFilter policy—create separate maps for each list. For detailed usage information, see Configuring Class Maps for Zone-Based Firewall Policies, page 21-19, Local Web Filter Class Add or Edit Match Criterion Dialog Boxes, page 21-29, and Configuring URLF Glob Parameter Maps, page 21-45.

SmartFilter (N2H2) or Websense Filtering—The class maps for N2H2 and Websense define any server response as the matching criterion. For detailed usage information, see Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Trend Micro Filtering – The Trend class map lets you select various Productivity Categories and Security Ratings, as defined by Trend Micro, that you want to target. For detailed usage information, see Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Besides the maps used to define content filtering, you can also configure the following maps for content filter rules:

- Inspect Parameters maps – Zone-based firewall inspection includes several general settings, all of which have default values that are appropriate for most networks. If you want to adjust any of these settings, you can create an Inspect Parameters map. In the Policy Object Manager, select **Maps > Parameter Maps > Inspect > Inspect Parameters**, and review the detailed usage information in Configuring Inspect Parameter Maps, page 21-31.
- HTTP policy map – If you want to use deep inspection on the individual HTTP packets in addition to Web filtering, you can configure an HTTP policy map by clicking **Configure** next to the Protocol field in the Action section of the Adding and Editing Zone-based Firewall Rules, page 21-62. The HTTP policy map incorporates HTTP class maps that define the type of traffic you want to match and then defines the action to take. For example, you can target traffic that includes Java applets. In the Policy Object Manager, select **Maps > Policy Maps > Inspect > HTTP (Zone Based IOS)**, and review the detailed usage information in Configuring Policy Maps for Zone-Based Firewall Policies, page 21-34, HTTP (IOS) Class Add or Edit Match Criterion Dialog Boxes, page 21-22, and Configuring Class Maps for Zone-Based Firewall Policies, page 21-19.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Zone-based Firewall Rules Page, page 21-58
- Creating Policy Objects, page 6-9
- Understanding Map Objects, page 6-78

Configuring Local Web Filter Parameter Maps

Use the Add and Edit Local Parameter Map dialog boxes to define a parameter map for local web filtering for zone-based firewall policies on routers. If you configure the action of a zone-based firewall policy rule as Content Filter, you can select a Web Filter policy map that incorporates a Local web filter parameter map (when you select Local for the parameter type on the Parameter tab). For more information about Web Filter policy maps, see Configuring Web Filter Maps, page 21-47.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Web Filter > Local** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-13 Add or Edit Local Web Filter Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Enable Alert	Whether to generate stateful packet inspection alert messages on the console.
Enable Allow Mode	Whether to allow or block URL requests when the URL filtering process does not have connectivity to a URL filtering database. When allow-mode is on, all unmatched URL requests are allowed; when off, all unmatched URL requests are blocked.
Block Page	The web page you want to present to the user if the user attempts to access a page that you block. You can select from the following: • None—The user is not presented with any information. • Message—The user is presented with the text message you enter in the edit box. • Redirect URL—The user is redirected to the URL you enter in the edit box.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring N2H2 or WebSense Parameter Maps

Use the Add and Edit N2H2 or Websense Parameter Map dialog boxes to define a parameter map for Smartfilter (N2H2) or Websense web filtering for zone-based firewall policies on routers. If you configure the action of a zone-based firewall policy rule as Content Filter, you can select a Web Filter policy map that incorporates an N2H2 or Websense web filter parameter map (when you select N2H2 or Websense for the parameter type on the Parameter tab). For more information about Web Filter policy maps, see Configuring Web Filter Maps, page 21-47.

Navigation Path

Select **Manage > Policy Objects**, then select N2H2 or WebSense from the **Maps > Parameter Maps > Web Filter** folder in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78

- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-14 Add or Edit N2H2 or WebSense Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
URL Filtering Server Table	The list of URL filtering servers and their attributes. • To add servers, click the Add button and fill in the Add External Filter dialog box (see Add or Edit External Filter Dialog Box, page 21-41). • To edit a server, select it and click the Edit button. • To delete a server, select it and click the Delete button.
Enable Alert	Whether to generate stateful packet inspection alert messages on the console.
Enable Allow Mode	Whether to allow or block URL requests when the URL filtering process does not have connectivity to a URL filtering database. When allow-mode is on, all unmatched URL requests are allowed; when off, all unmatched URL requests are blocked.
Block Page	The web page you want to present to the user if the user attempts to access a page that you block. You can select from the following: • None—The user is not presented with any information. • Message—The user is presented with the text message you enter in the edit box. • Redirect URL—The user is redirected to the URL you enter in the edit box.
Source Interface	The interface whose IP address should be used as the source IP address when a TCP connection is established between the system and the URL filtering server.
Maximum Cache Entries	The maximum number of entries to store in the categorization cache. The default is 5000.
Cache Life Time	How long, in hours, an entry remains in the cache table. The default is 24.
Maximum Requests	The maximum number of pending requests. The range is from 1 to 2147483647. The default is 1000.
Maximum Responses	The maximum number of HTTP responses that can be buffered. The range is from 0 and 20000. The default is 200.

Table 21-14 Add or Edit N2H2 or WebSense Parameter Map Dialog Boxes (continued)

Element	Description
Truncate Hostname Truncate Script Parameters	Whether to truncate the URLs: • If you do not select an option, URLs are not truncated. • If you select Hostname, URLs are truncated at the end of the domain name. • If you select Script Parameters, URLs are truncated at the left-most question mark in the URL. **Tip** Although you can select both options, it is illogical to do so.
Enable Server Log	Whether to send information about HTTP requests to the URL filtering server's log server. The information includes the URL, the hostname, the source IP address, and the destination IP address.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit External Filter Dialog Box

Use the Add or Edit External Filter dialog box to add a URL filtering server to an N2H2, Websense, or URL Filter parameter map policy object.

Navigation Path

Click the **Add** button beneath the server table, or select a server and click the **Edit** button, from any of the following dialog boxes:

- Add or Edit N2H2 or WebSense Parameter Map dialog boxes. See Configuring N2H2 or WebSense Parameter Maps, page 21-39.
- Add or Edit URL Filter Parameter Map dialog boxes. See Configuring URL Filter Parameter Maps, page 21-43.

Field Reference

Table 21-15 Add or Edit External Filter Dialog Box

Element	Description
Server	The fully-qualified domain name or IP address of the URL filtering server.
Port	The port that is listening for requests.
Retransmission Count	The number of times the router retransmits the lookup request when a response is not received from the server. The range is from 1 to 10.

Table 21-15 Add or Edit External Filter Dialog Box (continued)

Element	Description
Timeout	The number of seconds that the router waits for a response from the server. The range is from 1 to 300.
Outside	Whether the server is outside the network.

Configuring Trend Parameter Maps

Use the Add and Edit Trend Parameter Map dialog boxes to define a parameter map for Trend Micro web filtering for zone-based firewall policies on routers. If you configure the action of a zone-based firewall policy rule as Content Filter, you can select a Web Filter policy map that incorporates a Trend web filter parameter map (when you select Trend for the parameter type on the Parameter tab). For more information about Web Filter policy maps, see Configuring Web Filter Maps, page 21-47.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Web Filter > Trend** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-16 Add or Edit Trend Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Enable Allow Mode	Whether to allow or block URL requests when the URL filtering process does not have connectivity to a URL filtering database. When allow-mode is on, all unmatched URL requests are allowed; when off, all unmatched URL requests are blocked.
Block Page	The web page you want to present to the user if the user attempts to access a page that you block. You can select from the following: • None—The user is not presented with any information. • Message—The user is presented with the text message you enter in the edit box. • Redirect URL—The user is redirected to the URL you enter in the edit box.
Maximum Requests	The maximum number of pending requests. The range is from 1 to 2147483647. The default is 1000.

Table 21-16 Add or Edit Trend Parameter Map Dialog Boxes (continued)

Element	Description
Maximum Responses	The maximum number of HTTP responses that can be buffered. The range is from 0 and 20000. The default is 200.
Truncate Hostname	Whether to truncate URLs at the end of the domain name.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring URL Filter Parameter Maps

Use the Add and Edit URL Filter Parameter Map dialog boxes to define the parameters and match criterion and values for an inspection map used in a zone-based firewall policy for a router.

If you configure the action of a zone-based firewall policy rule as Content Filter, you can select a URL Filter parameter map to define web filtering parameters and match criteria. However, if the router is running Cisco IOS Software release 12.4(20)T or higher, the recommended approach is to configure a Web Filter policy map along with parameter and class maps for the appropriate server type (local, N2H2, Trend, or Websense). For more information, see Configuring Web Filter Maps, page 21-47.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Web Filter > URL Filter** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-17 Add or Edit URL Filter Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Local Filtering Tab The fields on this tab define the properties for local URL filtering.	

Table 21-17 Add or Edit URL Filter Parameter Map Dialog Boxes (continued)

Element	Description
Whitelisted and Blacklisted Domains tables	These tables define the domain names for which the software will not contact the external URL filtering server. Domain names on the allow list are always allowed. Domain names on the block list are always blocked. Use these lists to identify entire domains that you want to allow without restriction (such as your company's web site) or block completely (such as pornography sites). Domain names can be complete (including the host name, such as www.cisco.com), or partial (such as cisco.com). For partial names, all web site hosts on that domain are either permitted or denied. You can also enter host IP addresses. • To add a domain name, click the Add button and fill in the Add Server dialog box (see Add or Edit URL Domain Name Dialog Box for URL Filter Parameters, page 21-45). • To edit a domain name, select it and click the Edit button. • To delete a domain name, select it and click the Delete button.
Enable Alert	Whether to generate stateful packet inspection alert messages on the console.
Enable Audit Trail	Whether to log URL information to the syslog server or router.
Enable Allow Mode	Whether to allow or block URL requests when the URL filtering process does not have connectivity to a URL filtering database. When allow-mode is on, all unmatched URL requests are allowed; when off, all unmatched URL requests are blocked.
External Filtering Tab	
The fields on this tab define the properties for an external URL filtering server.	
Server Type Server Table	The type of external URL filtering server you are configuring, either SmartFilter (N2H2) or Websense. • To add servers, click the Add button and fill in the Add External Filter dialog box (see Add or Edit External Filter Dialog Box, page 21-41). • To edit a server, select it and click the Edit button. • To delete a server, select it and click the Delete button.
Source Interface	The interface whose IP address should be used as the source IP address when a TCP connection is established between the system and the URL filtering server.
Maximum Cache Entries	The maximum number of entries to store in the categorization cache. The default is 5000.
Maximum Requests	The maximum number of pending requests. The range is from 1 to 2147483647. The default is 1000.
Maximum Responses	The maximum number of HTTP responses that can be buffered. The range is from 0 and 20000. The default is 200.

Table 21-17 Add or Edit URL Filter Parameter Map Dialog Boxes (continued)

Element	Description
Truncate Hostname Truncate Script Parameters	Whether to truncate the URLs: • If you do not select an option, URLs are not truncated. • If you select Hostname, URLs are truncated at the end of the domain name. • If you select Script Parameters, URLs are truncated at the left-most question mark in the URL. Do not select any truncate options for devices running software releases lower than 12.4(15)T or you will receive a validation error. **Tip** Although you can select both options, it is illogical to do so.
Enable Server Log	Whether to send information about HTTP requests to the URL filtering server's log server. The information includes the URL, the hostname, the source IP address, and the destination IP address.
Additional Fields	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit URL Domain Name Dialog Box for URL Filter Parameters

Use the Add URL Domain Name dialog box to add web site domain names to the allow lists (allowed) or block lists (not allowed).

Domain names can be complete (including the host name, such as www.cisco.com), or partial (such as cisco.com). For partial names, all web site hosts on that domain are either permitted or denied. You can also enter host IP addresses.

Navigation Path

From the Add or Edit URL Filter Parameter Map dialog boxes, click the **Add** button beneath the whitelist or blacklist tables, or select a name and click the **Edit** button. See Configuring URL Filter Parameter Maps, page 21-43.

Configuring URLF Glob Parameter Maps

Use the Add and Edit URLF Glob Parameter Map dialog boxes to define a parameter map for the inspection of URLs in a Local web filter class map.

A single URLF Glob should contain only segments of URLs that you want to block or allow. Your goal is to create class maps ofallowed or blocked URLs. You can then define Local web filter policy maps to allow or block the identified URLs.

A single URLF Glob must also be limited to one of these types of URL segments:

- Strings that appear in the server name of a URL, which includes the name of the server and the domain name of the network. For example, www.cisco.com.
- Strings that appear in URL keywords, which are the strings that appear between / characters in a URL, or which are the file names. For example, in the URL segment www.cisco.com/en/US/, both en and US are keywords. The file name in a URL, such as index.html, is also considered a keyword.

You cannot use the characters /, {, }, and ? in a URLF glob.

To match a server name or URL keyword, the string in the URL must match exactly the string included in the URLF glob unless you use wildcard metacharacters to specify a variable string pattern. You can use the following metacharacters for pattern matching for either server names or URL keywords:

- * (Asterisk). Matches any sequence of zero or more characters. For example, *.edu matches all servers in the education domain, and you could use hack* to block www.example.com/hacksite/123.html.
- [abc] (Character class). Matches any character in the brackets. The character matching is case sensitive. For example, [abc] matches a, b, or c, but not A, B, or C. Thus, you could use www.[ey]xample.com to block both www.example.com and www.yxample.com.
- [a-c] (Character range class). Matches any character in the range. The character matching is case sensitive. [a-z] matches any lowercase letter. You can mix characters and ranges; for example, [abcq-z] matches a, b, c, q, r, s, t, u, v, w, x, y, z, and so does [a-cq-z].The dash (-) character is literal only if it is the last or the first character within the brackets, [abc-] or [-abc].
- [0-9] (Numerical range class). Matches any number in the brackets. For example [0-9] matches 0, 1, 2, 3, 4, 5, 6, 7, 8, or 9. Thus, you can use www.example[0-9][0-9].com to block www.example01.com, www.example33.com, and www.example99.com (and so forth).

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Parameter Maps > Web Filter > URLF Glob Parameters** in the table of contents. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Local Web Filter Class Add or Edit Match Criterion Dialog Boxes, page 21-29
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-18 Add or Edit URLF Glob Parameter Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.

Table 21-18 Add or Edit URLF Glob Parameter Map Dialog Boxes (continued)

Element	Description

Value	The server domains or keywords for the URLs you are targeting. Enter only one type of glob: either all server domains, or all URL keywords, but not a mixture of both.
	If you include more than one entry, separate the entries with new lines. For example, the following entries identify all government or education web servers:
	`*.gov`
	`*.edu`

Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring Web Filter Maps

Use the Add and Edit Web Filter Map dialog boxes to define the parameters and match criterion and values for an inspection map used in a zone-based firewall policy for a router.

If you configure the action of a zone-based firewall policy rule as Content Filter, you can select a Web Filter policy map to define web filtering parameters and match criteria. You can select Web Filter policy maps only for routers running Cisco IOS Software release 12.4(20)T and higher. If you are configuring zone-based firewalls for routers running Cisco IOS Software release 12.4(6)T up to 12.4(20)T, you must configure a URL Filter parameter map instead of a Web Filter policy map. For more information, see Configuring URL Filter Parameter Maps, page 21-43.

You can configure a mix of local and server-based web filtering. To do this, you should select a parameter map appropriate for the type of server you are using, and for the match criteria, an appropriate mix of local and server class maps. Do not mix class and parameter maps for different types of servers.

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > Policy Maps > Web Filter > Web Filter** from the Object Type selector. Right-click inside the table and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Map Objects, page 6-78
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-19 Add and Edit FTP Map Dialog Boxes

Element	Description
Name	The name of the policy object. A maximum of 40 characters is allowed.
Description	A description of the policy object. A maximum of 200 characters is allowed.
Parameters tab	
Parameter Type Parameter Map	The type of parameter map to include in the Web Filter policy map. Select None if you do not want to select a parameter map. If you select a specific parameter type, enter the name of the parameter map in the **Parameter Map** field. Click **Select** to select the map from a list or to create a new parameter map object.
Match Condition and Action Tab The Match All table lists class maps included in the policy map, and the action to take for traffic that matches the class. For traffic to match this class, all criteria defined in the selected class maps must be met. • To add a criterion, click the Add button and fill in the Match Condition and Action dialog box (see Add or Edit Match Condition and Action Dialog Boxes for Zone-Based Firewall and Web Filter Policies, page 21-35). • To edit a criterion, select it and click the Edit button. • To delete a criterion, select it and click the Delete button.	
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Changing the Default Drop Behavior

By default, all traffic between zones is dropped unless explicitly allowed. However, you can change this default behavior, as described in this section.

Security Manager converts the parameters—including class, parameter, and policy maps—that you supply for zone-based firewall rules into a series of IOS commands that the router will recognize. These are the so-called "CLI" (command line interface) configuration commands, which you can preview in separate window by choosing Tools > Preview Configuration. See Previewing Configurations, page 8-44 for more information. In addition, the section, Troubleshooting Zone-based Rules and Configurations, page 21-54, discusses an example of zone-based firewall CLI commands.

For the purposes of this discussion, the most interesting of these commands is **policy-map**, which is used to apply your zone policy for each pair of zones. That is, for any given zone-pair, all rules defining traffic (classes) and actions are applied within one policy-map. Further, Security Manager appends the current **class-default** class to the end of each policy-map's class list to capture any packets not processed by a zone rule.

The default class-default is drop—appending this class to each policy-map is how the implicit dropping of traffic between zones is accomplished. However, as mentioned, you can change this default behavior for any zone-pair. For example, you might elect to pass all unmatched traffic, or you might change the default to drop and log so you can determine what traffic is not being matched by your existing rules.

Note The only options for the default behavior are Drop, Drop and Log, Pass, and Pass and Log.

If you want the default policy to continue to drop packets, you do not have to do anything in Security Manager. This rule is generated automatically. If you do want to change the default behavior for a zone-pair, you must provide a **Permit any any IP** rule (that is, Match: Permit; Sources: any; Destinations: any; Services: IP in the Adding and Editing Zone-based Firewall Rules, page 21-62), with **Drop and Log**, **Pass**, or **Pass and Log** as the chosen Action. You must also ensure that this rule appears last in the list of rules for a zone pair. Security Manager interprets this as the intended class-default rule.

If your zone-based rules table includes a large number of rules, it might be difficult to ensure that this rule comes after all other rules for a zone pair. Here are a couple of techniques you can use to alleviate this:

- Use sections to organize the table, with one section per zone-pair. This will make it easier for you to order the rules for a zone-pair, as well as ensuring that the class-default rule comes last. For more information on working with sections, see Using Sections to Organize Rules Tables, page 12-20.
- Create a shared zone-based rules policy that includes the class-default rule in the Default scope, and inherit this rule in the device's local zone-based rule policy. For more information on inheritance and creating shared policies, see Inheriting or Uninheriting Rules, page 5-47 and Creating a New Shared Policy, page 5-54.

Configuring Settings for Zone-based Firewall Rules

Use the Zone Based Firewall settings page to: identify unreferenced zones; specify a zone for VPN interfaces; enable or disable WAAS support; maintain Trend Micro server and certificate information; and specify global Log settings on supported ASR devices.

Related Topics

- Zone Based Firewall Page, page 21-50
- Understanding the Zone-based Firewall Rules, page 21-3

Step 1 Access the Zone Based Firewall Page, page 21-50 as follows:

- (Device view) Select an IOS device and then select **Firewall > Settings > Zone Based Firewall** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Zone Based Firewall** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 (Optional) On the Zones tab: add, edit and delete unreferenced zones.

The Zones tab lists all unreferenced zones defined on the device; that is, zones without any associated interfaces, rules or policies. Unreferenced zones are usually found and listed during device discovery, but you also can create named, “empty” zones here.

Step 3 (Optional) On the VPN tab, supply the name of the zone specifically set up for VPN traffic.

This zone ensures that dynamic VPN traffic can be processed by the zone-based firewall rules on this router. See Using VPNs with Zone-based Firewall Policies, page 21-6 for more information.

Step 4 (Optional) On the WAAS tab, select Enable WAAS to enable Wide Area Application Services interoperability.

If this option is not enabled, packets being optimized by a WAAS device may be dropped because WAAS increases the TCP packet sequence number during the TCP handshake. This behavior may be viewed as a possible attack by the IOS device.

Step 5 (Optional) On the Content Filter Settings tab, provide server settings for Trend Micro-based content filtering.

To use Trend Micro-based content filtering, you must configure contact information for the Trend Micro server on this tab of the Zone Based Firewall page. This tab also provides links to Trend Micro registration and certificate download. You must have an active subscription with Trend Micro to utilize this form of content filtering, and you must download and install a valid subscription certificate on this IOS device.

For more information, see Zone Based Firewall Page - Content Filter Tab, page 21-52.

Step 6 (Optional) On the Global Parameters (ASR) tab, you can configure global, logging-related settings specific to ASR devices:

- Log Dropped Packets – Select this option to log all packets dropped by the device; syslog logging must be enabled to view the information.
- Log Flow export timeout rate – NetFlow logs are created after a flow either expires or is timed out, and it is important to put a time limit on how long a flow can be active before expiring. This value is maximum number of minutes a flow can remain active before it is expired. The value can be any integer from 1 to 3600; the default is 30.
- Log Flow export destination IP – The IP address or host name of the NetFlow collector to which flow data is to be sent.
- Log Flow export destination port – The UDP port monitored by the NetFlow collector for flow data.

Zone Based Firewall Page

Use the Zone Based Firewall page to configure and identify unreferenced zones, specify a VPN zone, enable or disable WAAS support, maintain Trend Micro server and certificate information, and specify global Log settings on supported ASR devices.

The following tabs are described in the table on this page:

- **Zones**
- **VPN**
- **WAAS**
- **Global Parameters (ASR)**

The **Content Filtering** tab is detailed in Zone Based Firewall Page - Content Filter Tab, page 21-52.

Navigation Path

To access the Zone Based Firewall page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Zone Based Firewall** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Zone Based Firewall** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and choose **Edit Firewall Settings > Zone Based Firewall**.

Related Topics

- Configuring Settings for Zone-based Firewall Rules, page 21-49
- Understanding the Zone-based Firewall Rules, page 21-3
- Adding Zone-Based Firewall Rules, page 21-13

Field Reference

Table 21-20 Zone Based Firewall Page

Element	Description
Zones tab	This tab displays the Zones table, which lists unreferenced zones; that is zones without any associated interfaces, rules or policies. Unreferenced zones are usually found and listed during device discovery, but you also can create named, “empty” zones here. The Zones table lists the following information for each unreferenced zone: • Zone – The name of the Zone/Interface Role. • Content – Any interfaces assigned to the zone. • Description – Any user-provided comments about the zone. To add a zone to this table, click the Add Row button and provide a Zone name in the Zone dialog box.
VPN tab	This tab presents the VPN Zone field; a zone entry in this field ensures that dynamic VPN traffic can be processed by the zone-based firewall rules on this router. See Using VPNs with Zone-based Firewall Policies, page 21-6 for more information about this zone. Enter or Select the zone through which VPN traffic will pass.
WAAS tab	This tab presents the Enable WAAS check box. Select this option to enable Wide Area Application Services interoperability. If this option is not enabled, packets being optimized by a WAAS device may be dropped because WAAS increases the TCP packet sequence number during the TCP handshake. This behavior may be viewed as a possible attack by the IOS device.
Content Filtering tab	This tab displays server settings and certificate links for Trend Micro-based content filtering. For more information, see Zone Based Firewall Page - Content Filter Tab, page 21-52.

Table 21-20 Zone Based Firewall Page (continued)

Element	Description
Global Parameters (ASR) tab	This tab displays global, logging-related settings specific to ASR devices. Configure these settings as follows: • Log Dropped Packets – Select this option to log all packets dropped by the device; syslog logging must be enabled to view the information. • Log Flow export timeout rate – NetFlow logs are created after a flow either expires or is timed out, and it is important to put a time limit on how long a flow can be active before expiring. This value is maximum number of minutes a flow can remain active before it is expired. The value can be any integer from 1 to 3600; the default is 30. • Log Flow export destination IP – The IP address or host name of the NetFlow collector to which flow data is to be sent. • Log Flow export destination port – The UDP port monitored by the NetFlow collector for flow data.

Zone Based Firewall Page - Content Filter Tab

To use Trend Micro-based content filtering, you must configure contact information for the Trend Micro server on this tab of the Zone Based Firewall page. This tab also provides links to Trend Micro registration and certificate download. You must have an active subscription with Trend Micro to utilize this form of content filtering, and you must download and install a valid subscription certificate on this IOS device.

Navigation Path

To access the Zone Based Firewall page, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Zone Based Firewall** from the Device selector.
- (Policy view) Select **Firewall > Settings > Zone Based Firewall** from the Policy selector.
- (Map view) Right-click a device and choose **Edit Firewall Settings > Zone Based Firewall**.

Related Topics

- Zone-based Firewall Rules Page, page 21-58
- Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36
- Understanding the Zone-based Firewall Rules, page 21-3
- Adding Zone-Based Firewall Rules, page 21-13

Field Reference

Table 21-21 Zone Based Firewall Page - Content Filter Tab

Element	Description
Trend Micro Server Settings	
Cache-entry-lifetime (hrs)	How long, in hours, a look-up request to the Trend Micro server remains in the router's local URL cache table. The allowed range is 0 to 120; the default value is 24.
Cache-size (KBytes)	The maximum amount of memory to be used by the router's local URL cache. The allowed range is 0 to 120,000 KB; the default value is 250.
Server	The fully-qualified domain name or IP address of the Trend Micro URL filtering server.
HTTP Port	The port the Trend Micro server is listening to for HTTP requests. The default is 80.
HTTPS Port	The port the Trend Micro server is listening to for HTTPS requests. The default is 443.
Retransmission Count	The number of times the router retransmits a look-up request when a response is not received from the server. The range is 1 to 10.
Retransmission Timeout	The number of seconds that the router waits for a response from the server. The range is 1 to 300.
Alert	Whether stateful packet inspection messages are copied to the syslog.
Trend Micro Server Certificate Download Links	
Link to download certificates	Opens the page for installing Trusted Authority Certificates on Cisco IOS Routers for Trend URL Filtering Support.
Link for product registration	Opens the page for Product License Registration. You must enter the Product Authorization Key and register the router.

Zone Dialog Box

Use the Add and Edit Zone dialog boxes to add and edit unreferenced zones—zones without any associated interfaces, rules or policies.

Navigation Path

To access the Add and Edit Zone dialog boxes, do one of the following:

- (Device view) Select a device, then select **Firewall > Settings > Zone Based Firewall** from the Device selector. Right-click inside the Zones table, then select **Add Row**, or right-click a line item, then select **Edit Row**.
- (Policy view) Select **Firewall > Settings > Zone Based Firewall** from the Policy selector. Right-click inside the table, then select **Add Row**, or right-click a line item, then select **Edit Row**.
- (Map view) Right-click a device and select **Edit Firewall Policies > Settings > Zone Based Firewall Rules**.

Enter a zone name in the Zone field, or click **Select** to choose one from the Interfaces Selector dialog box.

Related Topics

- Zone Based Firewall Page, page 21-50
- Understanding the Zone-based Firewall Rules, page 21-3
- Configuring Settings for Zone-based Firewall Rules, page 21-49

Troubleshooting Zone-based Rules and Configurations

Zone-based firewall rules are powerful, but also complex. Using zone rules, you can replace access rules, inspection rules, and Web filter rules with a single type of firewall rule. Because zone-based firewall rules can perform so many possible actions, the configuration generated from them uses many different types of configuration commands, including structures for access control lists (ACLs), class maps, and policy maps. There is no one-to-one correspondence between a zone-based firewall rule and a line in the configuration (unlike access rules, for example).

To illustrate this complexity, this topic describes the relationship between zone-based firewall rules and the configuration generated from them. You do not need to know any of the information in this topic to create and deploy zone-based firewall rules. However, if you are familiar with the CLI (command line interface), or if you find that your rules are generating undesired results, this information can help you understand and troubleshoot zone-based firewall rules.

Consider the set of rules shown in the following illustration. These rules form a policy for a single zone pair, affecting traffic moving from the Inside zone to the Outside zone. This is traffic from your internal network going to the Internet. The rules define the following actions:

- Drop all traffic from the 10.100.10.0/24 and 10.100.11.0/24 networks.
- Drop all FTP and FTPS traffic from the 10.100.12.0/24 network.
- Drop all peer-to-peer traffic from any network.
- Inspect (and allow) all FTP/FTPS traffic (except for that from 10.100.12.0/24, which is already dropped).
- Inspect all HTTP traffic using an additional deep-inspection policy map.
- And finally, perform generic inspection of all remaining TCP/UDP traffic.

Figure 21-3 Example of Zone-based Rules for a Zone Pair

No.	Permit	Source	Destination	Service	From Zone	To Zone	Inspected Protocol	Action
Local - Mandatory (7 Rules)								
1	✓	10.100.10.0/24	any	IP	Inside	Outside		Drop
2	✓	10.100.11.0/24	any	IP	Inside	Outside		Drop
3	✓	10.100.12.0/24	any	IP	Inside	Outside	Ftp Ftps	Drop
4	✓	any	any	IP	Inside	Outside	Bittorrent Edonkey Fasttrack Icq Kazaa2	Drop
5	✓	any	any	IP	Inside	Outside	Ftp Ftps	Inspect
6	✓	any	any	IP	Inside	Outside	Http(HTTPpmap)	Inspect
7	✓	any	any	IP	Inside	Outside	Tcp Udp	Inspect

194853

When you deploy these rules, Security Manager generates the following configuration. The bold letters are added for reference in the explanation that follows the configuration.

A.

```
class-map type inspect http match-any HTTPcmap
 match req-resp protocol-violation
 match request port-misuse any
!
```

B.

```
policy-map type inspect http HTTPpmap
 class type inspect http HTTPcmap
  reset
  log
!
```

C.

```
class-map type inspect CSM_ZBF_CLASS_MAP_1
 match access-group name CSM_ZBF_CMAP_ACL_1
!
```

D.

```
class-map type inspect match-any CSM_ZBF_CMAP_PLMAP_1
 match protocol ftp
 match protocol ftps
!
```

E.

```
class-map type inspect CSM_ZBF_CLASS_MAP_2
 match access-group name CSM_ZBF_CMAP_ACL_2
 match class-map CSM_ZBF_CMAP_PLMAP_1
!
```

F.

```
class-map type inspect match-any CSM_ZBF_CLASS_MAP_3
 match protocol bittorrent
 match protocol edonkey
 match protocol fasttrack
 match protocol icq
 match protocol kazaa2
!
```

G.

```
class-map type inspect CSM_ZBF_CLASS_MAP_4
 match protocol http
!
```

H.

```
class-map type inspect match-any CSM_ZBF_CLASS_MAP_5
 match protocol tcp
 match protocol udp
!
```

I.

```
policy-map type inspect CSM_ZBF_POLICY_MAP_1
 class type inspect CSM_ZBF_CLASS_MAP_1
  drop
 class type inspect CSM_ZBF_CLASS_MAP_2
  drop
 class type inspect CSM_ZBF_CLASS_MAP_3
  drop
class type inspect CSM_ZBF_CMAP_PLMAP_1
  inspect
 class type inspect CSM_ZBF_CLASS_MAP_4
  inspect
  service-policy http HTTPpmap
 class type inspect CSM_ZBF_CLASS_MAP_5
  inspect
 class class-default
  drop
!
```

J.

```
zone security Inside
zone security Outside
zone-pair security CSM_Inside-Outside_1 source Inside destination Outside
 service-policy type inspect CSM_ZBF_POLICY_MAP_1
!
interface GigabitEthernet0/1
 ip address dhcp
 zone-member security Inside
!
interface GigabitEthernet0/2
 ip address dhcp
 zone-member security Outside
!
```

K.

```
ip access-list extended CSM_ZBF_CMAP_ACL_1
 permit ip 10.100.10.0 0.0.0.255 any
 permit ip 10.100.11.0 0.0.0.255 any
!
```

L.

```
ip access-list extended CSM_ZBF_CMAP_ACL_2
 permit ip 10.100.12.0 0.0.0.255 any
```

!

The following list explains how the rules in Security Manager are converted to device-configuration commands, to aid your understanding of the relationship between the two. The list numbering corresponds to the rule numbers from the rules table in Security Manager (see the previous illustration):

1. This rule drops all traffic from the 10.100.10.0/24 network. The Permit, Source, Destination, and Service fields are used to create the first access control entry (ACE) in the ACL named CSM_ZBF_CMAP_ACL_1 defined in (K). This ACL is referenced from the class map CSM_ZBF_CLASS_MAP_1 defined in (C), which then defines the first drop rule in the policy map CSM_ZBF_POLICY_MAP_1, defined in (I).

   The policy map (I) is used to define the zone service policy in (J). Because this policy map is how all of the rules are assigned to the zone pair, (J) is not mentioned again.

2. This rule drops all traffic from the 10.100.11.0/24 network. This rule is combined with rule 1 by adding an ACE to the ACL defined in (K). The rest of the configuration is identical to rule 1. Thus, rules 1 and 2 essentially become a single rule in the device configuration.

3. This rule drops all FTP/FTPS traffic from the 10.100.10.12/24 network. The Permit, Source, Destination, and Service fields are used to create the ACL named CSM_ZBF_CMAP_ACL_2 defined in (L). The Protocol table generates the class map CSM_ZBF_CMAP_PLMAP_1 defined in (D), which specifies the FTP and FTPS protocols. The ACL and FTP/FTPS class map are then used in a new class map, CSM_ZBF_CLASS_MAP_2 defined in (E), which completes the characterization of the traffic based on the combination of source and protocol. Finally, (E) is referenced in the policy map (I) as the second rule.

4. This rule drops peer-to-peer traffic from any source that uses any of these protocols: Bittorrent, eDonkey, FastTrack, ICQ, or Kazaa2. This rule prevents any of your internal servers from being used as a file-sharing source for these services. Because the rule applies to all sources and destinations for the default IP service, no ACL is required. Instead, the configuration starts with the class map CSM_ZBF_CLASS_MAP_3 defined in (F). This class map is referenced by the third drop rule in the policy map (I).

5. This rule inspects FTP/FTPS traffic from any source to any destination, which means these services are allowed. Note that rule 3, because it comes above rule 5, already drops FTP/FTPS traffic from the 10.100.12.0/24 network, so the combination of these rules means that FTP/FTPS traffic is inspected for all sources except 10.100.12.0/24. Because the Protocol table specifies the same protocols as it does for rule 3, no new class map is needed. Instead, the policy map (I) simply refers to the class map (D) as the fourth class type, but this time with the Inspect action.

6. This rule inspects HTTP traffic and applies a deep-inspection policy map named HTTPpmap. The HTTPpmap policy map (B) defines the action to take when traffic matches the criteria defined in the class map HTTPcmap (A). These maps specify that any HTTP connection that violates the HTTP protocol, or that misuses ports, should be reset (dropped) and a syslog entry generated. (Protocol violation and port misuse can characterize Denial of Service attacks.) The combination of (A) and (B) define the deep-inspection rules for this policy.

   An additional class map, CSM_ZBF_CLASS_MAP_4, is needed to specify the HTTP protocol (G). Then, the fifth class type rule in the policy map (I) refers to class map (G) for inspection, and the service-policy command refers to the policy map (B) for deep inspection.

7. This rule provides generic inspection on TCP/UDP traffic, allowing and inspecting the remaining TCP/UDP traffic from the internal network to the Internet and back. The class map CSM_ZBF_CLASS_MAP_5 defined in (H) is generated from the Protocols table. This class map then becomes the next-to-last rule in the policy map (I).

8. Finally, there is an automatic rule, which appears as the final class-default rule in the policy map (I). This rule drops any traffic that does not match one of the class maps referenced in the policy map (I). For example, ICMP traffic from the internal network to the Internet will not be allowed. For information on configuring a different class-default rule, see Changing the Default Drop Behavior, page 21-48.

Zone-based Firewall Rules Page

Zone-based firewall rules provide unidirectional application of firewall policies between groups of interfaces known as "zones." That is, interfaces are assigned to zones, and specific inspection policies are applied to traffic moving between zones in one direction or the other.

A zone defines a boundary where traffic is subjected to specific restrictions as it crosses into another region of your network. The default zone-based firewall policy between zones is **deny all**. Thus, if no policy is explicitly configured, all traffic between zones is blocked.

Note Zone-based firewall policies can be configured only on Cisco IOS and ASR devices.

The Zone Based Firewall Rules page displays a list of currently configured zone-based firewall rules, and lets you add, edit and delete rules.

Tip Disabled rules are shown with hash marks covering the table row. When you deploy the configuration, disabled rules are removed from the device. For more information, see Enabling and Disabling Rules, page 12-20.

Navigation Path

To access the Zone Based Firewall Rules page, do one of the following:

- (Device view) Select a device, then select **Firewall > Zone Based Firewall Rules** from the Policy selector.
- (Policy view) Select **Firewall > Zone Based Firewall Rules** from the Policy Type selector. Create a new policy or select an existing one.
- (Map view) Right-click a device and select **Edit Firewall Policies > Zone Based Firewall Rules**.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Adding Zone-Based Firewall Rules, page 21-13
- Filtering Tables, page 1-48

Field Reference

Table 21-22 Zone Based Firewall Rules Page

Element	Description
No.	This number indicates the rule's position in the ordering of the list. You can use the Up Row and Down Row buttons to change the position of the selected rule.

Table 21-22 Zone Based Firewall Rules Page (continued)

Element	Description
Permit	Indicates whether the rule permits or denies traffic. • Permit – Shown as a green check mark. • Deny – Shown as a red circle with a slash.
Sources	The sources of traffic for this rule; can be networks or security groups. Multiple entries are displayed on separate lines within the table cell. • Network – The network, host, or IP address objects and definitions that are defined as the sources for this rule. The "All-Address" objects do not restrict the rule to specific hosts or networks. See Understanding Networks/Hosts Objects, page 6-80 and Specifying IP Addresses During Policy Definition, page 6-87 for additional information about these definitions. • Security Groups (IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+ only) – The names or tag numbers of the security groups defined as the sources for the rule. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. **Note** Security groups are part of the Cisco TrustSec Support for IOS feature which is supported on Cisco Integrated Services Router Generation 2 (ISR G2) only. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.

Table 21-22 Zone Based Firewall Rules Page (continued)

Element	Description
Destinations	The destinations of traffic for this rule; can be networks or security groups. Multiple entries are displayed on separate lines within the table cell. • Network – The network, host, or IP address objects and definitions that are defined as the destinations for this rule. The "All-Address" objects do not restrict the rule to specific hosts or networks. See Understanding Networks/Hosts Objects, page 6-80 and Specifying IP Addresses During Policy Definition, page 6-87 for additional information about these definitions. • Security Groups (IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+ only) – The names or tag numbers of the security groups defined as the destinations for the rule. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. **Note** Security groups are part of the Cisco TrustSec Support for IOS feature which is supported on Cisco Integrated Services Router Generation 2 (ISR G2) only. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.
Service	The services that define the types of traffic matched by this rule. Services are defined by objects that specify protocol and port information. See Understanding and Specifying Services and Service and Port List Objects, page 6-100 for more information.
From Zone	This rule applies only to traffic originating from this zone.
To Zone	This rule applies only to traffic destined for this zone.
Inspected Protocol	The protocol(s) on which the rule performs the chosen Action.

Table 21-22 Zone Based Firewall Rules Page (continued)

Element	Description
Action	Identifies how matched protocols are processed: • Drop – Matched traffic is silently dropped. The default action for all traffic. • Drop and Log – Matched traffic is logged and dropped. • Pass – The router forwards matched traffic from the source zone to the destination zone. • Pass and Log – Traffic is logged and forwarded. • Inspect – State-based traffic control; Inspect can provide application inspection and control for certain protocols, based on Port to Application Mapping (PAM). • Content Filter – HTTP content inspection based on a WebFilter parameter map, or a WebFilter policy map. **Note** The Log options generate system-log messages; you must ensure that syslog logging is configured to capture these messages.
Options	The Inspect Parameter map assigned to this rule; available only with Inspect and Content Filter actions.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of this rule, if provided. A maximum of 1024 characters is allowed.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Query button	To run policy queries, which can help you evaluate your rules and identify ineffective rules that you can delete. See Generating Policy Query Reports, page 12-28.
Find and Replace button (binoculars icon)	Searches for values in rules tables, such as IP addresses and policy object names, to facilitate locating and making changes to rules in tables. See Finding and Replacing Items in Rules Tables, page 12-16.
Up button	Moves the selected rule up one row in the table.
Down button	Moves the selected rule down one row in the table.
Add button	Opens the Add Zone-based Firewall Rule dialog box, where you can create a new rule.
Edit button	Used to edit the selected rule in the table; opens the Edit Zone-based Firewall Rule dialog box.
Delete button	Deletes the selected rule from the table.

Adding and Editing Zone-based Firewall Rules

Use the Add and Edit Zone based Firewall Rule dialog boxes to add and edit zone-based firewall rules on Cisco IOS and ASR devices.

Navigation Path

From the Zone-based Firewall Rules Page, page 21-58, click the **Add Row** button, or select a row and click the **Edit Row** button.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Configuring Settings for Zone-based Firewall Rules, page 21-49
- Adding Zone-Based Firewall Rules, page 21-13

Field Reference

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes

Element	Description
Enable Rule	When selected, the rule is enabled on the device after the configuration is generated and deployed. Deselect this option to disable the rule without deleting it.
Traffic	Define the traffic flow to which this rule is applied.
Match	Choose whether to Permit or Deny matched traffic. See Understanding the Relationship Between Permit/Deny and Action in Zone-based Firewall Rules, page 21-8 for additional information about this option.

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes (continued)

Element	Description
Sources	Provide traffic sources for this rule; can be networks or security groups. You can enter values, enter object names, or select objects for one or more of the following types of sources: **Note** Enter more than one value in any of these fields by separating the items with commas. • Network – You can specify various network, host, and IP address definitions, either individually or as objects. The "All-Address" objects do not restrict the rule to specific hosts or networks. See Understanding Networks/Hosts Objects, page 6-80 and Specifying IP Addresses During Policy Definition, page 6-87 for additional information about these definitions. • Security Groups (IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+ only) – Enter or Select the name or tag number for one or more source security groups for the rule, if any. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. **Note** Security groups are part of the Cisco TrustSec Support for IOS feature which is supported on Cisco Integrated Services Router Generation 2 (ISR G2) only. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes (continued)

Element	Description
Destinations	Provide traffic destinations for this rule; can be networks or security groups. You can enter values, enter object names, or select objects for one or more of the following types of sources: **Note** Enter more than one value in any of these fields by separating the items with commas. • Network – You can specify various network, host, and IP address definitions, either individually or as objects. The "All-Address" objects do not restrict the rule to specific hosts or networks. See Understanding Networks/Hosts Objects, page 6-80 and Specifying IP Addresses During Policy Definition, page 6-87 for additional information about these definitions. • Security Groups (IOS 15.2(2)T+ and IOS-XE 3.5.x(15.2(1)S)+ only) – Enter or Select the name or tag number for one or more destination security groups for the rule, if any. See Selecting Security Groups in Policies, page 14-16, Configuring TrustSec-Based Firewall Rules, page 14-17 and Creating Security Group Objects, page 14-14 for more information about security groups. **Note** Security groups are part of the Cisco TrustSec Support for IOS feature which is supported on Cisco Integrated Services Router Generation 2 (ISR G2) only. Each specification is combined with any others to limit traffic matches to only those flows that include all definitions. For example, specified user traffic originating from within a specified source address range.
Services	Specify the services that define the type of traffic to matched by this rule. You can enter any combination of service objects and service types (which are typically a protocol and port combination), separated by commas. See Understanding the Relationship Between Services and Protocols in Zone-based Firewall Rules, page 21-11 for additional information about this option. If you type in a service, you are prompted as you type with valid values. You also can click Select to select services from a list. For complete information on how to specify services, see Understanding and Specifying Services and Service and Port List Objects, page 6-100.
From Zone To Zone	Basic zone-based firewall rules are unidirectional; that is, they define a traffic flow that moves in only one direction between two zones. Enter or select the zone from which traffic flows can originate for this rule, and enter or select the zone to which traffic can flow.
Advanced button	Opens the Advanced Options dialog box where you can select time-range options. See Zone-based Firewall Rule: Advanced Options Dialog Box, page 21-67.

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes (continued)

Element	Description
Action	The action applied to traffic that matches this rule. Choose the desired Action:
Action: Drop, Drop and Log, Pass, Pass and Log	• **Drop** – Silently drops all packets for the specified Services. The default action for all traffic. • **Drop and Log** – Matched traffic is logged and dropped. • **Pass** – The router forwards matched packets from the From Zone to the To Zone. Return traffic is not recognized, so you have to specify additional rules for return traffic. This option is useful only for protocols such as IPsec-encrypted traffic. • **Pass and Log** – Traffic is logged and forwarded. For any of these Actions, you can select one or more protocols to be matched by clicking the **Select** button next to the Protocol table to open the Protocol Selector Dialog Box, page 21-68. However, this is not necessary; you can leave the Protocol table empty and pass or drop traffic based on the Sources, Destinations, and Services parameters; in effect, these are standard access rules. The Protocol Selector dialog box also provides access to the Configure Protocol Dialog Box, page 21-69, where you can edit the Port Application Mapping (PAM) parameters for the selected protocol. **Note** The Log options generate system-log messages; you must ensure that syslog logging is configured to capture these messages.

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes (continued)

Element	Description
Action: Inspect	**Inspect** provides state-based traffic control—the device maintains connection or session information for TCP and UDP traffic, meaning return traffic in reply to connection requests is permitted. Choose this option to apply packet inspection based on your selected Layer 4 (TCP, UDP) and Layer 7 (HTTP, IMAP, instant messaging, and peer-to-peer) protocols. You also can edit PAM settings for the selected protocols, and you can set up deep packet inspection (DPI) and provide additional protocol-related information for the Layer 7 protocols. See Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16 for more information. 1. You can select one or more protocols for inspection by clicking the **Select** button next to the Protocol table to open the Protocol Selector Dialog Box, page 21-68. 2. The Protocol Selector dialog box also provides access to the Configure Protocol Dialog Box, page 21-69, where you can create custom protocols, and edit the PAM and DPI parameters for the selected protocol. 3. **Inspect Parameters** – You can apply a customized set of connection, timeout, and other settings by entering the name of an Inspect Parameter map in this field, or you can click Select to select one from a list. You also can create new Inspect Parameter maps from the selection-list dialog box; see Configuring Inspect Parameter Maps, page 21-31 for more information. If you do not specify an Inspect Parameters map, the default settings are used.

Table 21-23 Add and Edit Zone based Firewall Rule Dialog Boxes (continued)

Element	Description
Action: Content Filter	**Content Filter** provides URL filtering based on a supplied parameter or policy map. The router intercepts HTTP requests, performs protocol-related inspection, and optionally contacts a third-party server to determine whether the requests should be allowed or blocked. You can provide a WebFilter parameter map, which defines filtering based on local URL lists, as well as information from an external SmartFilter (previously N2H2) or Websense server. Alternately, you can provide a WebFilter policy map that accesses Local, N2H2, Websense, or Trend Micro filtering data. 1. When **Content Filter** is the chosen Action, HTTP is the specified Protocol. You can click Configure to open the Configure Protocol Dialog Box, page 21-69, where you can edit the HTTP PAM settings, and apply an HTTP DPI map. 2. Select **WebFilter Parameter Map**, or **WebFilter Policy Map**, and supply the name of an appropriate map. You can click the appropriate Select button to select the map from a list; you also can create new maps from the selection-list dialog box. See Configuring Content Filtering Maps for Zone-based Firewall Policies, page 21-36 for information about configuring these maps. 3. **Inspect Parameters** – You can apply a customized set of connection, timeout, and other settings by entering the name of an Inspect Parameter map in this field, or you can click Select to select one from a list. You also can create new Inspect Parameter maps from the selection-list dialog box; see Configuring Inspect Parameter Maps, page 21-31 for more information. If you do not specify an Inspect Parameters map, the default settings are used.
Description	(Optional) You can enter a description of up to 1024 characters to help you identify the rule when viewing the rules table.
Category	(Optional) You can assign a category to the rule, to help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Zone-based Firewall Rule: Advanced Options Dialog Box

Use the Zone-Based Firewall Rule Advanced Options dialog box to apply specific time-range information to a zone-based firewall rule.

Navigation Path

In the Traffic section of the Add or Edit Zone based Firewall Rule dialog box, click the **Advanced** button.

Related Topics

- Adding and Editing Zone-based Firewall Rules, page 21-62
- Understanding the Zone-based Firewall Rules, page 21-3

Field Reference

Table 21-24 Advanced Options Dialog Box

Element	Description
Time Range	This feature lets you define time periods during which this zone-based firewall rule is active. If you do not specify a time range, the rule is immediately and always active. Enter the name of a time-range object, or click **Select** to choose one from a list in the Time Ranges Selector dialog box. You can create and edit time-range objects from this dialog box. For more information, see Configuring Time Range Objects, page 6-71.
Options	This feature lets you apply an initial-packet-fragment or an established-connection restriction to this zone-based firewall rule. Choose one of the following options: • **None**—No packet-fragment or established-connection restrictions are applied. • **Fragment** – If chosen, the rule is applied to non-initial packet fragments; the fragment is either permitted or denied accordingly. The white paper, "Access Control Lists and IP Fragments," provides additional information that is also relevant to zone-based firewall rules. • **Established** – For the TCP protocol only; requires an established connection. A match occurs if the TCP datagram has the ACK or RST control bits set. The non-matching case is that of the initial TCP datagram to form a connection.

Protocol Selector Dialog Box

Use the Protocol Selector dialog box to specify one or more communication protocols as part of the definition of traffic for a zone-based firewall rule.

The Protocol Selector dialog box also provides access to the Configure Protocol dialog box, which you can use to create custom protocols and edit Port Application Mapping (PAM) parameters for existing protocols. The Configure Protocol dialog box is also where you select Deep Inspection policy maps, and Protocol Info parameter maps, for certain protocols. See Configure Protocol Dialog Box, page 21-69 for more information.

Navigation Path

The Protocol Selector dialog box can be accessed from the Add and Edit Zone based Firewall Rule dialog boxes (described in Adding and Editing Zone-based Firewall Rules, page 21-62). In either dialog box, choose any Action except Content Filter and then click the Select button next to the Protocol table.

You can also open the Protocol Selector dialog box by right-clicking the Inspected Protocol column for any entry in the Zone Based Firewall Rules table, and then choosing Edit Protocols.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Adding and Editing Zone-based Firewall Rules, page 21-62

- Selecting Objects for Policies, page 6-2
- Configure Protocol Dialog Box, page 21-69

Table 21-25 Protocol Selector Dialog Box

Element	Description
Available Protocols	A list of protocols that can be selected for a zone-based firewall rule. **Tip** You can create a custom protocol by clicking the Create button below the Selected Protocols column, opening the Configure Protocol Dialog Box, page 21-69.
Selected Protocols	The list of protocols you have selected for this zone-based firewall rule. **Tip** You can edit Port Application Mapping (PAM) settings for the protocol highlighted in the Selected Protocols column: click the Edit button below the Selected Protocols column to open the Configure Protocol Dialog Box, page 21-69.
>> button	Moves the highlighted protocols from the Available Protocols column to the Selected Protocols column. You can select multiple protocols using the standard Shift-click and Ctrl+click functions.
<< button	Moves the highlighted protocols from the Selected Protocols column back to the Available Protocols column. You can select multiple protocols using the standard Shift-click and Ctrl+click functions.

Configure Protocol Dialog Box

Packet inspection can be configured in zone-based firewall rules by the selection of specific protocol objects, which define Port Application Mapping (PAM) parameters (Layer 4 protocols and ports, and optionally specific networks and hosts). A Layer 7 (HTTP, IMAP, Instant Messaging, and peer-to-peer) protocol can also include a deep-packet inspection policy specific to that protocol. Refer to Adding and Editing Zone-based Firewall Rules, page 21-62 for information about selecting protocols during zone-based firewall rule definition.

The Configure Protocol dialog box is used to edit existing protocol definitions, and to create custom definitions, for use with zone-based firewall rules. For example, if a protocol does not use its default ports for some or all networks, you can configure different port mappings.

Navigation Path

The Configure Protocol dialog box is accessed from the Protocol Selector Dialog Box, page 21-68, as follows:

- Click the Create (+) button below the Selected Protocols list to create a new protocol.
- Select a protocol in the Selected Protocols list, and click the Edit (pencil) button to edit that protocol.

Related Topics

- Understanding the Zone-based Firewall Rules, page 21-3
- Adding Zone-Based Firewall Rules, page 21-13
- Protocol Selector Dialog Box, page 21-68

Table 21-26 Configure Protocol Dialog Box

Element	Description
Protocol Name	The name of the selected protocol. If you are creating a custom protocol, you can enter a name of up to 19 characters. Custom protocol names must begin with **user-**.
Enable Signature	This option is available only when editing the peer-to-peer (eDonkey, FastTrack, Gnutella, Kazaa2) protocols. Enabling this option means Network-based Application Recognition (NBAR) heuristics will be applied to the traffic to detect "telltales" that signify specific P2P application activity. These telltales includes port-hopping and other changes in application behavior to avoid traffic detection. **Note** This level of traffic inspection comes at the price of increased CPU utilization and reduced network throughput capability.
Deep Inspection	This option is available only when editing the H.323, HTTP, IM (AOL, ICQ, MSN Messenger, Windows Messenger, and Yahoo Messenger), IMAP, P2P (eDonkey, FastTrack, Gnutella, Kazaa2), POP3, SIP, SMTP, Sun RPC protocols, and **Inspect** is the chosen Action for the zone-based firewall rule. Enter or Select the name of the Inspect policy map to be used with the selected protocol. See Configuring Inspection Maps for Zone-based Firewall Policies, page 21-16 for more information about these policy maps.
Protocol Info	This option is available only when editing the Instant Messaging (AOL, ICQ, MSN Messenger, Windows Messenger, and Yahoo Messenger) and Stun-ice protocols. Enter or Select the name of the Protocol Info parameter map to be used with the selected protocol. These parameter maps define the DNS servers that interact with these applications, which helps the Instant Messaging (IM) application engine recognize the IM traffic and enforce the configured policy for that IM application. See Configuring Protocol Info Parameter Maps, page 21-33 for more information about these parameter maps.
Port Application Mapping	These options let you customize the Port Application Mapping (PAM) parameters for the selected protocol.
Protocol	Select the transport protocol(s) for this mapping: • TCP/UDP • TCP • UDP
Ports	Enter any combination of a single port number, multiple port numbers, or a range of ports (for example, 60000-60005). Separate multiple entries with commas. Do not specify a range that overlaps already mapped ports.

Table 21-26 Configure Protocol Dialog Box (continued)

Element	Description
Networks	If this protocol/port mapping is only for specific networks or hosts, enter the names or IP addresses of the networks or hosts, or the names of the network/host objects. You can click Select to open the Networks/Hosts Selector. Separate multiple entries with commas.

Managing Traffic Zones

You can assign multiple interfaces to a traffic zone, which lets traffic from an existing flow exit or enter the ASA on any interface within the zone. This capability allows Equal-Cost Multi-Path (ECMP) routing on the ASA as well as external load balancing of traffic to the ASA across multiple interfaces.

Non-Zoned Behavior

The Adaptive Security Algorithm takes into consideration the state of a packet when deciding to permit or deny the traffic. One of the enforced parameters for the flow is that traffic enters and exits the same interface. Any traffic for an existing flow that enters a different interface is dropped by the ASA.

Traffic zones let you group multiple interfaces together so that traffic entering or exiting any interface in the zone fulfills the Adaptive Security Algorithm security checks.

Related Topics

- Why Use Zones?, page 22-1
- ECMP Routing, page 22-4
- Understanding Traffic Zones, page 22-6
- Prerequisites for Traffic Zones, page 22-7
- Guidelines for Traffic Zones, page 22-8
- Configuring Traffic Zones, page 22-9

Why Use Zones?

Asymmetric Routing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to asymmetric routing on the destination network, return traffic arrived from ISP 2 on the Outside2 interface.

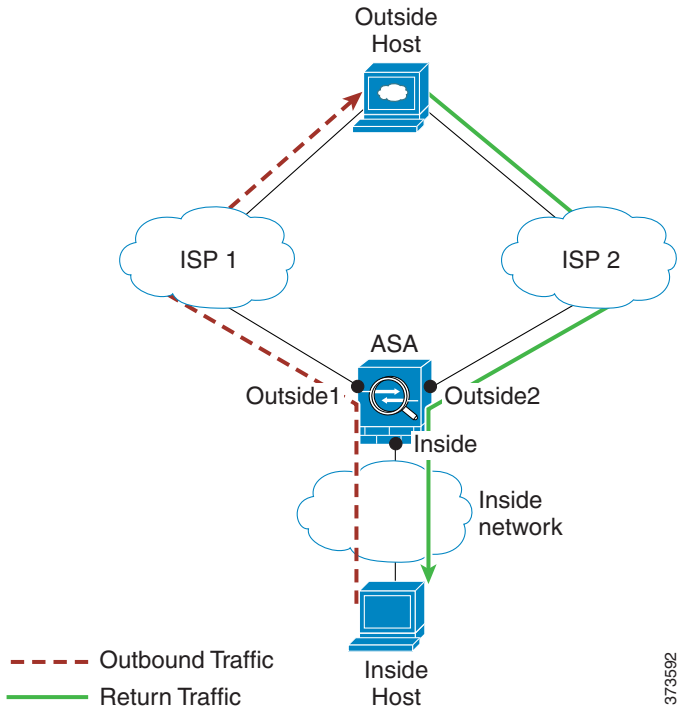

Non-Zoned Problem: The ASA maintains the connection tables on a per-interface basis. When the returning traffic arrives at Outside2, it will not match the connection table and will be dropped.

Zoned Solution: The ASA maintains connection tables on a per-zone basis. If you group Outside1 and Outside2 into a zone, then when the returning traffic arrives at Outside2, it will match the per-zone connection table, and the connection will be allowed.

Lost Route

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. Due to a lost or moved route between Outside1 and ISP 1, traffic needs to take a different route through ISP 2.

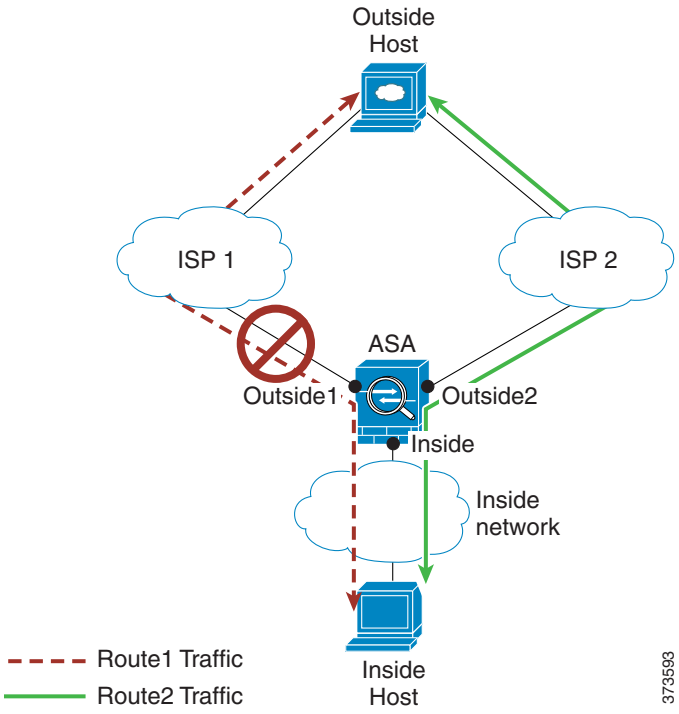

Non-Zoned Problem: The connection between the inside and outside host will be deleted; a new connection must be established using a new next-best route. For UDP, the new route will be used after a single packet drop, but for TCP, a new connection has to be reestablished.

Zoned Solution: The ASA detects the lost route and switches the flow to the new path through ISP 2. Traffic will be seamlessly forwarded without any packet drops.

Load Balancing

In the following scenario, a connection was established between an inside host and an outside host through ISP 1 on the Outside1 interface. A second connection was established through an equal cost route through ISP 2 on Outside2.

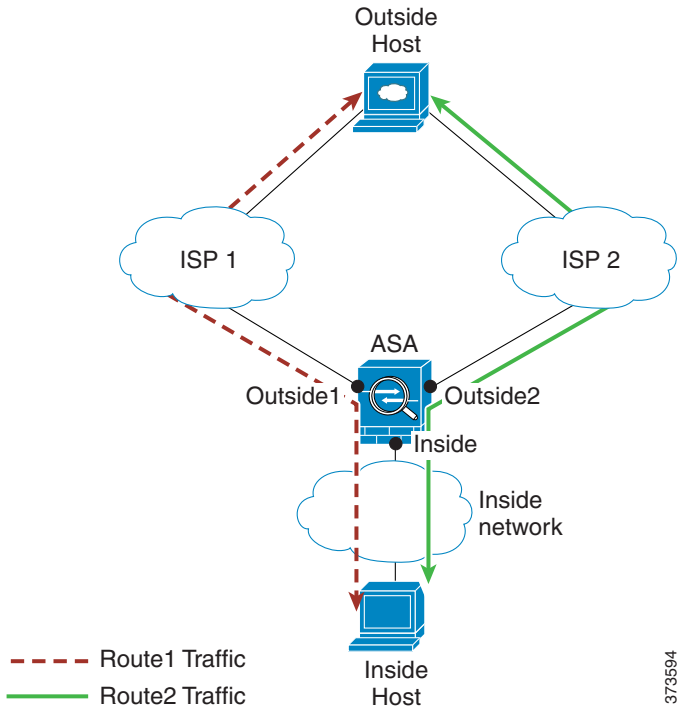

Non-Zoned Problem: Load-balancing across interfaces is not possible; you can only load-balance with equal cost routes on one interface.

Zoned Solution: The ASA load-balances connections across up to eight equal cost routes on all the interfaces in the zone.

Related Topics

- Chapter 22, "Managing Traffic Zones"
- ECMP Routing, page 22-4
- Understanding Traffic Zones, page 22-6
- Prerequisites for Traffic Zones, page 22-7
- Guidelines for Traffic Zones, page 22-8
- Configuring Traffic Zones, page 22-9

ECMP Routing

The ASA supports Equal-Cost Multi-Path (ECMP) routing.

Non-Zoned ECMP Support

Without zones, you can have up to three equal cost static or dynamic routes per interface. For example, you can configure three default routes on the outside interface that specify different gateways:

```
route outside 0 0 10.1.1.2
route outside 0 0 10.1.1.3
```

```
route outside 0 0 10.1.1.4
```

In this case, traffic is load-balanced on the outside interface between 10.1.1.2, 10.1.1.3, and 10.1.1.4. Traffic is distributed among the specified gateways based on an algorithm that hashes the source and destination IP addresses.

ECMP is not supported across multiple interfaces, so you cannot define a route to the same destination on a different interface. The following route is disallowed when configured with any of the routes above:

```
route outside2 0 0 10.2.1.1
```

Zoned ECMP Support

With zones, you can have up to 8 equal cost static or dynamic routes across up to 8 interfaces within a zone. For example, you can configure three default routes across three interfaces in the zone:

```
route outside1 0 0 10.1.1.2
route outside2 0 0 10.2.1.2
route outside3 0 0 10.3.1.2
```

Similarly, your dynamic routing protocol can automatically configure equal cost routes. The ASA load-balances traffic across the interfaces with a more robust load balancing mechanism.

When a route is lost, the ASA seamlessly moves the flow to a different route.

How Connections Are Load-Balanced

The ASA load balances connections across equal cost routes using a hash made from the packet 6-tuple (source and destination IP address, source and destination port, protocol, and ingress interface). Unless the route is lost, a connection will stay on the chosen interface for its duration.

Packets within a connection are not load-balanced across routes; a connection uses a single route unless that route is lost.

The ASA does not consider the interface bandwidth or other parameters when load balancing. You should make sure all interfaces within the same zone have the same characteristics such as MTU, bandwidth, and so on.

The load-balancing algorithm is not user configurable.

Falling Back to a Route in Another Zone

When a route is lost on an interface, if there are no other routes available within the zone, then the ASA will use a route from a different interface/zone. If this backup route is used, then you may experience packet drops as with non-zoned routing support.

Related Topics

- Chapter 22, “Managing Traffic Zones”
- Why Use Zones?, page 22-1
- Understanding Traffic Zones, page 22-6
- Prerequisites for Traffic Zones, page 22-7
- Guidelines for Traffic Zones, page 22-8
- Configuring Traffic Zones, page 22-9

Understanding Traffic Zones

Interface-Based Security Policy

Zones allow traffic to and from any interface in the zone, but the security policy itself (access rules, NAT, and so on) is still applied per interface, not per zone. If you configure the same security policy for all interfaces within the zone, then you can successfully implement ECMP and load balancing for that traffic. For more information about required parallel interface configuration, see Prerequisites for Traffic Zones, page 22-7.

Supported Services for Traffic Zones

The following services are supported with zones:

- Access Rules
- NAT
- Service Rules, except for QoS traffic policing.
- Routing

You can also configure to- and from-the-box services (see below), although full zoned support is not available.

Do not configure other services (such as VPN or Botnet Traffic Filter) for interfaces in a traffic zone; they may not function or scale as expected.

Note For detailed information about how to configure the security policy, see Prerequisites for Traffic Zones, page 22-7.

Security Levels

The first interface that you add to a zone determines the security level of the zone. All additional interfaces must have the same security level. To change the security level for interfaces in a zone, you must remove all but one interface, and then change the security levels, and re-add the interfaces.

Primary and Current Interface for the Flow

Each connection flow is built based on the initial ingress and egress interfaces. These interfaces are the *primary* interfaces.

If a new egress interface is used because of route changes or asymmetric routing, then the new interfaces are the *current* interfaces.

Joining or Leaving a Zone

When you assign an interface to a zone, any connections on that interface are deleted. The connections must be reestablished.

If you remove an interface from a zone, any connections that have the interface as the primary interface are deleted. The connections must be reestablished. If the interface is the current interface, the ASA moves the connections back to the primary interface. The zone route table is also refreshed.

To- and From-the-Box Traffic

- You cannot add management-only or management-access interfaces to a zone.
- For management traffic on regular interfaces in a zone, only asymmetric routing on existing flows is supported; there is no ECMP support.

- You can configure a management service on only one zone interface, but to take advantage of asymmetric routing support, you need to configure it on all interfaces. Even when the configurations are parallel on all interfaces, ECMP is not supported.
- The ASA supports the following to- and from-the-box services in a zone:
  - Telnet
  - SSH
  - HTTPS
  - SNMP
  - Syslog
  - BGP

Overlapping IP Addresses Within a Zone

For non-zoned interfaces, the ASA supports overlapping IP address networks on interfaces so long as you configure NAT properly. However, overlapping networks are not supported on interfaces in the same zone.

Related Topics

- Chapter 22, "Managing Traffic Zones"
- Why Use Zones?, page 22-1
- ECMP Routing, page 22-4
- Prerequisites for Traffic Zones, page 22-7
- Guidelines for Traffic Zones, page 22-8
- Configuring Traffic Zones, page 22-9

Prerequisites for Traffic Zones

- Configure all interface parameters including the name, IP address, and security level. Note that the security level must match for all interfaces in the zone. You should plan to group together like interfaces in terms of bandwidth and other Layer 2 properties.
- Configure the following services to match on all zone interfaces:
  - Access Rules—Apply the same access rule to all zone member interfaces, or use a global access rule.
  - NAT—Configure the same NAT policy on all member interfaces of the zone or use a global NAT rule.

    Interface PAT is not supported.

Note When you use interface-specific NAT and PAT pools, the ASA cannot switch connections over in case of the original interface failure.

If you use interface-specific PAT pools, multiple connections from the same host might load-balance to different interfaces and use different mapped IP addresses. Internet services that use multiple concurrent connections may not work correctly in this case.

– Service Rules—Use the global service policy, or assign the same policy to each interface in a zone.

  QoS traffic policing is not supported.

Note For VoIP inspections, zone load balancing can cause increased out-of-order packets. This situation can occur because later packets might reach the ASA before earlier packets that take a different path. Symptoms of out-of-order packets include:

–Higher memory utilization at intermediate nodes (firewall and IDS) and the receiving end nodes if queuing is used.

–Poor video or voice quality.

To mitigate these effects, we recommend that you use IP addresses only for load distribution for VoIP traffic.

- Configure routing with ECMP zone capabilities in mind.

Related Topics

- Chapter 22, "Managing Traffic Zones"
- Why Use Zones?, page 22-1
- ECMP Routing, page 22-4
- Understanding Traffic Zones, page 22-6
- Guidelines for Traffic Zones, page 22-8
- Configuring Traffic Zones, page 22-9

Guidelines for Traffic Zones

Firewall Mode

Supported in routed firewall mode only. Does not support transparent firewall mode.

Failover

- You cannot add the failover or state link to a zone.
- In Active/Active failover mode, you can assign an interface in each context to an asymmetrical routing (ASR) group. This service allows traffic returning on a similar interface on the peer unit to be restored to the original unit. You cannot configure both ASR groups and traffic zones within a context. If you configure a zone in a context, none of the context interfaces can be part of an ASR group.
- Only the primary interfaces for each connection are replicated to the standby unit; current interfaces are not replicated. If the standby unit becomes active, it will assign a new current interface if necessary.

Clustering

You cannot add the cluster control link to a zone.

Additional Guidelines

- You can create a maximum of 256 zones.
- You can only add physical interfaces to a zone.
- An interface can be a member of only one zone.
- You can include up to 8 interfaces per zone.
- The first interface that you add to a zone determines the security level of the zone. All additional interfaces must have the same security level.
- For ECMP, you can add up to 8 equal cost routes per zone, across all zone interfaces. You can also configure multiple routes on a single interface as part of the 8 route limit.

Related Topics

- Chapter 22, "Managing Traffic Zones"
- Why Use Zones?, page 22-1
- ECMP Routing, page 22-4
- Understanding Traffic Zones, page 22-6
- Prerequisites for Traffic Zones, page 22-7
- Configuring Traffic Zones, page 22-9

Configuring Traffic Zones

You can assign multiple interfaces to a traffic zone, which lets traffic from an existing flow exit or enter the ASA on any interface within the zone. This capability allows Equal-Cost Multi-Path (ECMP) routing on the ASA as well as external load balancing of traffic to the ASA across multiple interfaces.

Related Topics

- Understanding Interface Role Objects, page 6-73
- Chapter 22, "Managing Traffic Zones"
- Why Use Zones?, page 22-1
- ECMP Routing, page 22-4
- Understanding Traffic Zones, page 22-6
- Prerequisites for Traffic Zones, page 22-7
- Guidelines for Traffic Zones, page 22-8

Step 1 Do one of the following:

- (Device view) Select **Firewall > Settings > Zone** from the Policy selector.
- (Policy view) Select **Firewall > Settings > Zone** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Click the **Add Row** button beneath the Zone table to open the Zone dialog box.

Step 3 Enter the name of an interface role that identifies the interfaces that belong to the traffic zone you are configuring and then click **OK**. For more information about interface role object, see Understanding Interface Role Objects, page 6-73.

Tip You can click Select to select the interface role from a list of interface role objects or to define a new interface role object.

Step 4 Click **Save** to save the changes.

CHAPTER 23

Managing Transparent Firewall Rules

Transparent firewall rules are access control rules for non-IP layer 2 traffic. You can use these rules to permit or drop traffic based on the Ethertype value in the layer-2 packet.

This chapter contains the following topics:

- Configuring Transparent Firewall Rules, page 23-1
- Transparent Rules Page, page 23-3

Configuring Transparent Firewall Rules

Transparent firewall rules are access control rules for non-IP layer 2 traffic. You can use these rules to permit or drop traffic based on the Ethertype value in the layer-2 packet. These rules create Ethertype access control lists on the device. With transparent rules, you can control the flow of non-IP traffic across the device. (To control IP traffic, use access rules; see Understanding Access Rules, page 16-1.)

Transparent firewalls are devices that you place within a single subnet to control traffic flow across a bridge. They allow you to insert a firewall on a subnet without renumbering your networks.

You can configure transparent rules only on the following types of interfaces:

- **IOS 12.3(7)T or higher devices**—On layer-3 interfaces that are part of a bridge group:
  - Configure the interfaces you want to bridge as layer 3 in the **Interfaces > Interfaces** policy.
  - Configure a bridge group with two or more layer 3 interfaces in the **Platform > Device Admin > Bridging** policy (see Bridging on Cisco IOS Routers, page 63-18 and Defining Bridge Groups, page 63-20).
  - Create a bridge group virtual interface (BVI) using the same number as the bridge group (see Bridge-Group Virtual Interfaces, page 63-19). For example, if you create bridge group 12, create BVI12.
- **ASA, PIX 7.0+, FWSM devices**—On any interface when the device is running in transparent mode. If you are using multiple contexts, configure the rules on the individual security contexts.

  There are several other bridging policies that you can configure in the **Platform > Bridging** policy group including: ARP table and ARP inspection, MAC table and the ability to disable MAC learning, and the ability to configure a management IP address so that you can remotely manage the device. For more detail about transparent firewalls, see Chapter 47, "Configuring Bridging Policies on Firewall Devices" and Interfaces in Routed and Transparent Modes, page 46-5.

Tip On ASA, PIX, and FWSM in transparent mode, you must configure access rules to allow any IP traffic to pass through the device. Transparent rules control layer 2 non-IP traffic only.

Also, see NAT in Transparent Mode, page 24-16 for information about using network address translation on security devices.

You can also configure other types of firewall rules on these interfaces. The other types of rules apply to layer-3 and higher traffic.

Tip If you configure any transparent rule, an implicit **deny all** rule is added at the end of the rule list for each interface. You must ensure that you permit all desired traffic. You might want to include a **permit any** (for ASA/PIX/FWSM devices) or **permit 0x0000 0xFFFF** (for IOS devices) rule as the final rule in the table if your desire is simply to deny specific types of traffic, rather than permitting only specific types of traffic.

Related Topics

- Adding and Removing Rules, page 12-9
- Editing Rules, page 12-10
- Enabling and Disabling Rules, page 12-20

Step 1 Do one of the following to open the Transparent Rules Page, page 23-3:

- (Device view) Select **Firewall > Transparent Rules** from the Policy selector for a supported device type.
- (Policy view) Select **Firewall > Transparent Rules** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the row after which you want to create the rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add and Edit Transparent Firewall Rule Dialog Boxes, page 23-5.

Tip If you do not select a row, the new rule is added at the end of the local scope. You can also select an existing row and edit either the entire row or specific cells. For more information, see Editing Rules, page 12-10.

Step 3 Configure the rule. Following are the highlights of what you typically need to decide. For specific information on configuring the fields, see Add and Edit Transparent Firewall Rule Dialog Boxes, page 23-5.

- Permit or Deny—Whether you are allowing traffic that matches the rule or dropping it.
- Interfaces—The interface or interface role for which you are configuring the rule.
- The direction of traffic to which this rule should apply (in or out). The default is in.
- EtherType—The hexadecimal code or keyword (for ASA/PIX/FWSM only) that identifies the traffic. For a list of codes, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "Ether Type." For ASA/PIX/FWSM, you can select a keyword to identify some EtherTypes. For ASA/PIX/FWSM, the code must be 0x0600 at minimum.
- Mask—For rules applied to IOS devices, you must also specify a mask to apply to the EtherType. Use 0xFFFF to have the EtherType interpreted literally.

If you want to create a single rule to apply to a group of EtherTypes, convert the EtherTypes to binary and calculate an appropriate mask where 1 means to interpret the EtherType literally, and 0 means that any value should be allowed in the position. You must then convert your mask into hexadecimal.

Click **OK** when you are finished defining your rule.

Step 4 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.

Step 5 (IOS devices only) If you are configuring transparent rules on an IOS device, you can forward DHCP traffic across the bridge without inspection. To configure this, select the **Firewall > Settings > Inspection** policy and select the **Permit DHCP Passthrough (Transparent Firewall)** option. This setting is not supported on all IOS versions, so carefully inspect validation results to see if it will be configured on your device.

Transparent Rules Page

Use the Transparent Rules page to control access for non-IP layer-2 traffic. (To control IP traffic access, use access rules; see Understanding Access Rules, page 16-1.)

Transparent rules are limited to transparent firewalls, which are ASA, PIX 7.0+, and FWSM devices running in transparent mode, or layer-3 interfaces that are part of a bridge group on IOS 12.3(7)T+ devices. When deployed, transparent rules become Ethertype access control lists.

Configure the same rules on all bridged interfaces to allow traffic to pass both ways through the device.

For more detailed information about configuring transparent firewalls and the device requirements for deploying these rules, see Configuring Transparent Firewall Rules, page 23-1.

Tip Disabled rules are shown with hash marks covering the table row. When you deploy the configuration, disabled rules are removed from the device. For more information, see Enabling and Disabling Rules, page 12-20.

Navigation Path

To access Transparent Rules, do one of the following:

- (Device view) Select **Firewall > Transparent Rules** from the Policy selector for a supported device type.
- (Policy view) Select **Firewall > Transparent Rules** from the Policy Type selector. Select an existing policy or create a new one.
- (Map view) Right-click a device and select **Edit Firewall Policies > Transparent Rules**.

Related Topics

- Interfaces in Routed and Transparent Modes, page 46-5.
- Chapter 47, “Configuring Bridging Policies on Firewall Devices”
- Bridging on Cisco IOS Routers, page 63-18
- Defining Bridge Groups, page 63-20
- Bridge-Group Virtual Interfaces, page 63-19

- Filtering Tables, page 1-48

Field Reference

Table 23-1 Transparent Rules Page

Element	Description
No.	The ordered rule number.
Permit	Whether a rule permits or denies traffic based on the conditions set: • Permit—Shown as a green check mark. • Deny—Shown as a red circle with slash.
EtherType	The Ethernet packet type, which is the EtherType value in the packet. This can be a hexadecimal code or a keyword.
Mask	The 16-bit hexadecimal mask for the EtherType (for IOS devices only). A mask of 0xFFFF indicates the EtherType is literal. Any other mask indicates the corresponding bits in the EtherType to ignore. You must convert the hexadecimal number to binary to fully interpret the mask (binary 1 means interpret the corresponding EtherType value literally, 0 means allow any value at that position).
Interface	The interfaces or interface roles to which the rule is assigned. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. Multiple entries are displayed as separate subfields within the table cell. See Understanding Interface Role Objects, page 6-73.
Dir.	The direction of the traffic to which this rule applies: • In—Packets entering the interface. • Out—Packets exiting the interface.
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	The description of the rule, if any.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope or section. For more information, see Moving Rules and the Importance of Rule Order, page 12-19.
Add Row button	Click this button to add a rule to the table after the selected row using the Add and Edit Transparent Firewall Rule Dialog Boxes, page 23-5. If you do not select a row, the rule is added at the end of the local scope. For more information about adding rules, see Adding and Removing Rules, page 12-9.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells. For more information, see Editing Rules, page 12-10.

Table 23-1 Transparent Rules Page (continued)

Element	Description
Delete Row button	Click this button to delete the selected rule.

Add and Edit Transparent Firewall Rule Dialog Boxes

Use the Add and Edit Transparent Firewall Rule dialog boxes to add and edit transparent firewall rules, which are configured as EtherType access control lists on the device. Before you configure transparent rules, read Configuring Transparent Firewall Rules, page 23-1.

Navigation Path

From the Transparent Rules Page, page 23-3, click the **Add Row** button or select a row and click the **Edit Row** button.

Related Topics

- Interfaces in Routed and Transparent Modes, page 46-5.
- Chapter 47, "Configuring Bridging Policies on Firewall Devices"
- Bridging on Cisco IOS Routers, page 63-18
- Defining Bridge Groups, page 63-20
- Bridge-Group Virtual Interfaces, page 63-19
- Editing Rules, page 12-10
- Adding and Removing Rules, page 12-9

Field Reference

Table 23-2 Add and Edit Transparent Firewall Rule Dialog Boxes

Element	Description
Enable Rule	Whether to enable the rule, which means the rule becomes active when you deploy the configuration to the device. Disabled rules are shown overlain with hash marks in the rule table. For more information, see Enabling and Disabling Rules, page 12-20.
Action	Whether the rule permits or denies traffic based on the conditions you define.
Interfaces	The interfaces or interface roles to which the rule is assigned. You must select only bridged, transparent interfaces (for more specific information, see Configuring Transparent Firewall Rules, page 23-1). Enter the name of the interface or the interface role, or click **Select** to select the interface or role from a list, or to create a new role. An interface must already be defined to appear on the list. Interface role objects are replaced with the actual interface names when the configuration is generated for each device. See Understanding Interface Role Objects, page 6-73.

Table 23-2 Add and Edit Transparent Firewall Rule Dialog Boxes (continued)

Element	Description
Traffic Direction	The direction of the traffic to which this rule applies: • In—Packets entering an interface. • Out—Packets exiting an interface.
EtherType	The hexadecimal code or keyword (for ASA/PIX/FWSM only) that identifies the traffic based on the EtherType value in the packet. Enter or select the following: • The hexadecimal EtherType value. For a list of codes, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "Ether Type." – IOS devices—You can enter any value from 0x0000 to 0xFFFF. – ASA/PIX/FWSM devices—The value must be 0x0600 or higher. • For ASA/PIX/FWSM devices, you can also select these keywords: – bpdu—Spanning Tree Bridge Protocol Data Units – ipx—Internet Packet Exchange – mpls-unicast—Multi-Protocol Label Switching, unicast. – mpls-multicast—MPLS multicast. – isis—IS-IS pass-through – any—Any packet regardless of EtherType. – eii-ipx – raw-ipx **Tip** The keyword "isis" in the list above refers to IS-IS pass-through support, which is new in Security Manager 4.4. "IS-IS pass-through support" means that IS-IS traffic can flow through the ASA in transparent mode. **Note** Beginning from 4.16, the ethertype dsap CLI is used to interpret the installed ACEs—regardless of whether it was created with ether type bpdu, ipx, or isis— in ether type dsap format. This feature is supported for ASA 9.9(1) and higher devices.
Wildcard Mask (IOS)	The mask is a 16-bit hexadecimal number that determines how the EtherType code is interpreted. A mask of 0xFFFF indicates the EtherType is literal. Any other mask indicates the corresponding bits in the EtherType to ignore. You must convert the hexadecimal number to binary to fully interpret the mask (binary 1 means interpret the corresponding EtherType value literally, 0 means allow any value at that position).
Category	The category assigned to the rule. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Description	An optional description of the rule (up to 1024 characters).

Edit Transparent EtherType Dialog Box

Use the Edit Transparent EtherType dialog box to edit the EtherType in a transparent firewall rule. Enter the hexadecimal code that identifies the traffic. For ASA/PIX/FWSM devices, you can also select the keyword for some types of traffic. For a list of codes, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "Ether Type." For a more detailed description of EtherType, see Add and Edit Transparent Firewall Rule Dialog Boxes, page 23-5.

For more information, see Configuring Transparent Firewall Rules, page 23-1.

Navigation Path

Right-click the EtherType cell in a transparent rule (on the Transparent Rules Page, page 23-3) and select **Edit EtherType**. You can edit the EtherType for one row at a time.

Edit Transparent Mask Dialog Box

Use the Edit Transparent Mask dialog box to edit the mask in a transparent firewall rule for an IOS device. The mask is a 16-bit hexadecimal number that determines how the EtherType code is interpreted.

A mask of 0xFFFF indicates the EtherType is literal. Any other mask indicates the corresponding bits in the EtherType to ignore. You must convert the hexadecimal number to binary to fully interpret the mask (binary 1 means interpret the corresponding EtherType value literally, 0 means allow any value at that position).

For more information, see Configuring Transparent Firewall Rules, page 23-1.

Navigation Path

Right-click the Mask cell in a transparent rule (on the Transparent Rules Page, page 23-3) and select **Edit Mask**. You can edit the mask for one row at a time.

Configuring Network Address Translation

These topics provide conceptual information about network address translation (NAT) in general, and about translation types and various implementations:

- Understanding Network Address Translation, page 24-2
  - Types of Address Translation, page 24-3
  - About "Simplified" NAT on ASA 8.3+ Devices, page 24-4

The following topics describe configuration and management of NAT rules on various Cisco devices:

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS, it does not support any enhancements.

Cisco IOS routers

- NAT Policies on Cisco IOS Routers, page 24-5
  - NAT Page: Interface Specification, page 24-6
  - NAT Page: Static Rules, page 24-6
  - NAT Page: Dynamic Rules, page 24-10
  - NAT Page: Timeouts, page 24-13

PIX, FWSM, and ASA security devices

- NAT Policies on Security Devices, page 24-15
- NAT in Transparent Mode, page 24-16
- Translation Options Page, page 24-17
- **PIX, FWSM, and pre-8.3 ASA devices**
  - Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
  - Address Pools, page 24-18
  - Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- **ASA 8.3+ devices**
  - Configuring NAT on ASA 8.3+ Devices, page 24-33
  - Translation Rules: ASA 8.3+, page 24-34
  - Add and Edit NAT Rule Dialog Boxes, page 24-36

– Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42
– Per-Session NAT Rules: ASA 9.0(1)+
– Add and Edit Per Session NAT Rule Dialog Boxes, page 24-47

Understanding Network Address Translation

Address translation substitutes the real address in a packet with a mapped address that is routable on the destination network. As part of the process, the device also records the substitution in a translation database; these records are known as "xlate" entries. The appropriate xlate entry must exist to allow address translation on return packets—the substitution of the original real address for the mapped address; this procedure is sometimes referred to as "untranslation." Thus, network address translation (NAT) actually consists of two steps: the translation of a real address into a mapped address, and the reverse translation for returning traffic.

One of the main functions of NAT is to enable private IP networks to connect to the Internet. Network address translation replaces a private IP address with a public IP address, translating the private addresses in the internal network into legal, routable addresses that can be used on the public Internet. In this way, NAT conserves public addresses; for example, NAT rules can be configured to utilize only one public address for the entire network in communications with the outside world.

Other functions of NAT include:

- Security – Keeping internal IP addresses hidden discourages direct attacks.
- IP routing solutions – Overlapping IP addresses are not a problem.
- Flexibility – You can change internal IP addressing schemes without affecting the public addresses available externally. For example, for a server accessible to the Internet, you can maintain a fixed IP address for Internet use, but internally, you can change the server address.

Cisco devices support both NAT, which provides a globally unique address for each outbound host session, and Port Address Translation (PAT), which provides the same single address combined with a unique port number, for up to 64,000 simultaneous outbound or inbound host sessions. The global addresses used for NAT come from a pool of addresses specifically designated for address translation. The unique global address that is used for PAT can be either one global address, or the IP address of a given interface.

The device translates an address when an existing NAT rule matches the specific traffic. If no NAT rule matches, processing for the packet continues. The exception is when you enable NAT control. NAT control requires that packets traversing from a higher security interface (inside) to a lower security interface (outside) match a NAT rule, or processing for the packet stops.

Cisco devices can perform NAT or PAT on both inbound and outbound connections. This ability to translate inbound addresses is called "Outside NAT" because addresses on the outside, or less secure, interface are translated to a usable inside IP address. Just as when you translate outbound traffic, you may choose dynamic NAT, static NAT, dynamic PAT, or static PAT. If necessary, you can use outside NAT together with inside NAT to translate the both source and destination IP addresses of a packet.

Note In this document, all types of translation are generally referred to as NAT; see Types of Address Translation, page 24-3 for descriptions of the various types. When describing NAT, the terms inside and outside represent the security relationship between any two interfaces. The higher security level is inside and the lower security level is outside.

The release of ASA version 8.3 provides a simplified, interface-independent approach to configuring network address translation, as compared to earlier ASA versions and other devices. See About "Simplified" NAT on ASA 8.3+ Devices, page 24-4 for more information.

Related Topics

- Types of Address Translation, page 24-3
- About "Simplified" NAT on ASA 8.3+ Devices, page 24-4

Types of Address Translation

The following table briefly describes the various types of address translation.

Table 24-1 Types of Address Translation

Static NAT	Fixed translation of real source addresses to specific mapped addresses—each source address is always translated to the same mapped address, regardless of IP protocol and port number.
Static PAT	Fixed translation of real source addresses with specific TCP or UDP port numbers, to specific mapped addresses and ports. That is, each source address/port is always translated to the same mapped address/port.
Policy Static NAT	Fixed translation of real source addresses to specific mapped addresses. Destination networks/hosts are also specified, and the service is always IP.
Policy Static PAT	Fixed translation of real source addresses with specific TCP or UDP port numbers, to specific mapped addresses and ports. Destination networks/hosts and services are also specified.
Dynamic NAT	Dynamic translation of real source addresses to mapped addresses obtained from a pool of shared addresses. Each source address can be mapped to any available address in the pool.
Dynamic PAT	Translation of real source addresses to a single mapped address; singularity is provided by dynamic translation of related port numbers. That is, each real address/port combination is translated to the same mapped address, but assigned a unique port. This is sometimes referred to as "overloading."
Policy Dynamic NAT	Dynamic translation of specific source-address/destination-address/service combinations on a given interface, using a pool of shared addresses. Translation direction—outbound or inbound—is also specified.
Identity NAT	The specified address is translated to itself—that is, it is effectively not translated; applies to outbound connections only. Identity NAT is a particular type of Static NAT.
NAT Exempt	Translation is bypassed for specified source/destination address combinations; connections can be initiated in both the outbound and inbound directions.

Note While certain of these types do not apply to ASA 8.3 and later devices, the ASA 8.3+ devices do provide a Dynamic NAT and PAT option, which is Dynamic NAT with a Dynamic PAT back-up feature.

About "Simplified" NAT on ASA 8.3+ Devices

The release of ASA version 8.3 provides a simplified approach to configuring network address translation (NAT), as compared to earlier ASA versions and other devices. Configuration of NAT was simplified by replacing the earlier flow-based scheme with an "original packet" to "translated packet" approach.

All NAT rules on the device—static NAT, dynamic PAT, and dynamic NAT—are presented in a single table, and essentially the same dialog box is used to configure all NAT rules. The NAT rules are interface independent (that is, interfaces are optional), meaning the rules are independent of security levels also.

NAT rules are no longer dependent on security levels. A global address space consisting of all interfaces is available, and is specified using the keyword "any." All Interface fields default to **any**, so unless a specific interface is provided, the rule is applicable to all interfaces.

Network Object NAT

You also can define NAT properties on Host, Address Range, and Network objects, such that corresponding NAT rules are applied automatically to the designated security device. Using these objects means you need enter the necessary IP addresses, services, ports, and optional interfaces only once. These automatically generated, object-based rules are referred to as "Network Object NAT" rules. Note that these rules cannot be created or deleted from the rules table; you must edit the appropriate objects in the Policy Object Manager. You can, however, edit these rules from the rules table after they have been defined for the network object. For more information, see Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42.

Note Network Object NAT rules are not displayed in the Translation Rules table in Policy View because these rules are device-specific.

The NAT Table

As mentioned, all NAT rules on a device are presented in a single table, which is divided into three sections: a "manual" section, the Network Object NAT rules section, and another manual-rules section. You can add, edit and order rules in both manual sections; the Network Object NAT rules are added and ordered automatically, and as mentioned, to edit these rules you must edit the related objects.

The NAT rules in the table are applied on a top-down, first-match basis. That is, a packet is translated only when it matches a NAT rule, and as soon as a match is made, regardless of its location or section, NAT rule processing stops.

You can use this table to organize and manage the manual rules—you can insert rules in any order, and you can re-order them. The two sections of manual rules are provided to let you order manual rules both before and after the automatic object rules.

Network Object NAT rules are automatically ordered such that static rules appear before dynamic rules. These two types are each further ordered as follows:

- Fewest number of IP addresses – Rules for objects with one IP address are listed before those for objects with two addresses, which are before those with three addresses, and so on.
- IP address numbers – For objects having the same number of IP addresses, the rules are arranged such that the IP addresses themselves are in numerical order, from lower to higher. For example, 10.1.1.1 rules are listed before 11.1.1.1 rules.
- Object names – If the IP address is the same, the rules are ordered by alphabetizing the object names.

And remember, translation is based on the first matching rule.

Destination Translation

With manual static rules, in addition to source address translation, you also can configure destination address translation. Source and destination translation are defined at the same time, in the same dialog box. Again, while source translation can be static or dynamic, destination translation is always static, and is only available with manual rules.

Bi-directional or Twice NAT

When creating a manual static rule, you can select the "Bi-directional" option, which will produce an entry in the rules table that actually represents two static NAT rules, encompassing both translation directions. That is, a static rule is created for the specified source/translated address pairing, along with a mirror rule for the translated address/source pairing.

For example, if Bi-directional is chosen when you create a static rule with Host1 in the Source field and Host2 in the Translated field, two lines are added to the rules table: one with Host1 being translated to Host2, and one with Host2 being translated to Host1.

This is sometimes referred to as "Twice NAT" because only one look-up is required to fetch and process what is in effect two rules.

Many-to-one Addressing

Generally, static NAT rules are configured with one-to-one address mapping. However, you can now define static NAT rules in which many IP addresses map to a few or one IP address. Functionally, many-to-few is the same as many-to-one, but because the configuration is more complicated, we recommend creating a many-to-one rule for each address as needed.

Many-to-one addressing might be useful, for example, in a situation where a range of public IP addresses is used to reach a load balancer which redirects requests to an internal network.

Related Topics

- Configuring NAT on ASA 8.3+ Devices, page 24-33
- Add and Edit NAT Rule Dialog Boxes, page 24-36
- Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42

NAT Policies on Cisco IOS Routers

You can configure NAT policies on a Cisco IOS router from the following tabs on the NAT policy page:

- NAT Page: Interface Specification, page 24-6
- NAT Page: Static Rules, page 24-6
- NAT Page: Dynamic Rules, page 24-10
- NAT Page: Timeouts, page 24-13

Network Address Translation (NAT) converts private, internal LAN addresses into globally routable IP addresses. NAT enables a small number of public IP addresses to provide global connectivity for a large number of hosts.

For more information, see Understanding Network Address Translation, page 24-2.

Navigation Path

- (Device view) Select **NAT** from the Policy selector.

- (Policy view) Select **NAT (Router)** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

NAT Page: Interface Specification

Before creating NAT rules, you must define the "direction" of the traffic to be translated by specifying the Inside and Outside interfaces. Inside interfaces typically connect to a LAN that the router serves. Outside interfaces typically connect to your organization's WAN or to the Internet. You must designate at least one Inside interface and one Outside interface to enable the router to perform network address translation.

The Inside and Outside designations are used when interpreting translation rules: addresses connected to the Inside interface are translated to addresses on the Outside interface. After these interfaces are defined, they are used in all static and dynamic NAT translation rules.

Use the Interface Specification tab of the NAT policy page to specify the Inside and Outside interfaces.

Navigation Path

- (Device view) Select **NAT** from the Policy selector, then click the **Interface Specification** tab.
- (Policy view) Select **NAT (Router) > Translation Rules** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Interface Specification** tab.

Defining the Inside and Outside Interfaces

In the **NAT Inside Interfaces** and **NAT Outside Interfaces** fields, enter or Select the names of the interfaces or interface roles for the Inside and Outside interfaces, respectively. Separate multiple names or roles with commas (for example, Ethernet1/1, Ethernet1/2). Note that you cannot enter the same name in both fields.

Related Topics

- NAT Policies on Cisco IOS Routers, page 24-5
- NAT Page: Static Rules, page 24-6
- NAT Page: Dynamic Rules, page 24-10
- NAT Page: Timeouts, page 24-13
- NAT Page: Timeouts, page 24-13

NAT Page: Static Rules

You define a static NAT rule by specifying a local address that must be translated, as well as the global address to which it is translated. This is a static or fixed mapping—the local address is always translated to the same global address.

You can define static NAT rules that translate the addresses of single hosts, as well as static rules that translate multiple addresses in a subnet. When multiple local addresses must use the same global address, you must define the necessary port redirection information, which defines a different port for each local address using the global address.

Note We strongly recommend that you do not perform NAT on traffic that will be transmitted over a VPN. Translating addresses on this traffic causes it to be sent out unencrypted instead of encrypted over the VPN.

The procedure for creating a static rule depends on whether the address being translated represents a port, a single host, or an entire subnet:

- You define a **static NAT rule for a single host** by entering the original address to translate and the global address to which it is translated. The global address may be taken from an interface on the device.
- You define a **static NAT rule for a subnet** by entering one of the addresses in the subnet (including the subnet mask) as the original address, and one of the global addresses that you want to use as the translated address. The router configures the remaining addresses based on the subnet mask you provide.
- You define a **static NAT rule for a port** by entering the original IP address and the global address to which it should be translated. The global address may be taken from an interface on the device. In addition, you must select the protocol used by the port, as well as the local and global port numbers.

The Add Static NAT Rule and Edit Static NAT Rule dialog boxes are used to add and edit these rules. Refer to NAT Static Rule Dialog Boxes, page 24-7 for descriptions of the fields displayed in the table on this page.

Before You Begin

- Define the inside and outside interfaces used for NAT. See NAT Page: Interface Specification, page 24-6.

Navigation Path

- (Device view) Select **NAT** from the Policy selector, then click the **Static Rules** tab.
- (Policy view) Select **NAT (Router) > Translation Rules** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Static Rules** tab.

Related Topics

- NAT Policies on Cisco IOS Routers, page 24-5
- NAT Page: Dynamic Rules, page 24-10
- NAT Page: Timeouts, page 24-13
- Standard Security Manager rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

NAT Static Rule Dialog Boxes

Use the Add/Edit NAT Static Rule dialog boxes to add or edit static address translation rules. Except for their titles, the two dialog boxes are identical.

Navigation Path

Go to the NAT Page: Static Rules, page 24-6 tab; click the **Add** button beneath the table to add a new rule, or select a rule in the table and click **Edit** to update that rule.

Related Topics

- Understanding Interface Role Objects, page 6-73

Field Reference

Table 24-2 Add/Edit NAT Static Rule Dialog Boxes

Element	Description
Static Rule Type	The type of local address to be translated by this static rule: • Static Host – A single host requiring static address translation. • Static Network – A subnet requiring static address translation. • Static Port – A single port requiring static address translation. If you select this option, you must define the Port Redirection parameters.
Original Address	An IP address, or the name of a network/host object representing the address(es) to be translated. You can enter or Select the object name. Network/host objects are logical collections of IP addresses that represent networks, hosts, or both. See Understanding Networks/Hosts Objects, page 6-80 for more information. **Note** Do not enter a local address belonging to this router, as it could cause Security Manager management traffic to be translated. Translating this traffic will cause a loss of communication between the router and Security Manager.
Translated Address	Use the options in this section of the dialog box to specify the address(es) to which the Original Address(es) are translated: • **Specify IP** – Select this option to specify an IP address, or the name of a network/host object that provides the translated address(es). Add an IP address, or the name of a network/host object, in the **Translated IP/Network** field. You can enter or Select the object name. • **Use Interface IP** – Select this option to specify that the IP address assigned to a particular interface be used as the translated address. Enter or Select the name of the desired **Interface**. (This is typically the interface from which translated packets leave the router.) **Note** This option is not available when Static Network is the chosen rule type. Only one static rule may be defined per interface.

Table 24-2 Add/Edit NAT Static Rule Dialog Boxes (continued)

Element	Description
Port Redirection	These parameters specify port information for the address translations. Port address translation lets you to use the same public IP address for multiple devices as long as the port specified for each device is different. **Note** These parameters are available only when Static Port is the chosen rule type. **Redirect Port** – When Static Port is chosen as the rule type, this box is automatically checked; it cannot be changed. Enter the appropriate information in the following fields: • **Protocol** – The communications protocol used for these ports: TCP or UDP. • **Local Port** – The port number on the source network. Valid values range from 1 to 65535. • **Global Port** – The port number on the destination network that the router is to use for this translation. Valid values range from 1 to 65535.
Advanced	This section contains optional, advanced translation options. **Note** The Advanced options are available only when the Specify IP option is the selected method for defining the translated address(es). • No Alias – When selected, disables automatic aliasing for the global IP address translation. If the NAT pool used as an inside global pool consists of addresses on an attached subnet, an alias is generated for that address so that the router can answer Address Resolution Protocol (ARP) requests for those addresses. When deselected, global address aliases are permitted. • No Payload – When selected, prohibits an embedded address or port in the payload from being translated. The payload option performs NAT between devices on overlapping networks that share the same IP address. When an outside device sends a DNS query to reach an inside device, the local address inside the payload of the DNS reply is translated to a global address according to the relevant NAT rule. You can disable this feature by selecting the No Payload option. Otherwise, embedded addresses and ports in the payload may be translated. See Disabling the Payload Option for Overlapping Networks, page 24-10 for more information. • Create Extended Translation Entry – When checked, extended translation entries (addresses and ports) are created in the translation table. This lets you associate multiple global addresses with a single local address. This is the default. When this option is deselected, simple translation entries are created, allowing association of a single global address with the local address. **Note** This option is not available when Static Port is the chosen rule type.

Disabling the Payload Option for Overlapping Networks

Overlapping networks result when you assign an IP address to a device on your network that is already legally owned and assigned to a different device on the Internet or outside network. Overlapping networks can also result after the merger of two companies using RFC 1918 IP addresses in their networks. These two networks need to communicate, preferably without your having to re-address all their devices.

This communication is achieved as follows. The outside device cannot use the IP address of the inside device because it is the same as the address assigned to itself (the outside device). Instead, the outside device sends a Domain Name System (DNS) query for the inside device's domain name. The source of this query is the IP address of the outside device, which is translated to an address from a designated address pool. The DNS server located on the inside network replies with the IP address associated with the inside device's domain name in the data portion of the packet. The destination address of the reply packet is translated back to the outside device's address, and the address in the data portion of the reply packet is translated to an address from a different address pool. In this way, the outside device learns that the IP address for the inside device is one of the addresses from that second address pool, and it uses this address when it communicates with the inside device. The router running NAT takes care of the translations at this point.

To disable the translation of the address inside the payload, check the **No Payload** option when you create a static NAT rule based on a global IP translation.

NAT Page: Dynamic Rules

Use the NAT Dynamic Rules tab of the router's NAT page to manage dynamic address translation rules. A dynamic address translation rule dynamically maps hosts to addresses, using either the IP address of a specific interface (with dynamic port translation), or the addresses included in an address pool that are globally unique in the destination network.

Defining Dynamic NAT Rules

You define a dynamic NAT rule by first selecting an access control list (ACL) whose rules specify the traffic requiring translation.

Then, you must either select an interface with an IP address to which the addresses should be translated, or define a pool of addresses to be used. You define the pool by specifying a range of addresses and giving the range a unique name; you can specify multiple ranges. The router uses the available addresses in the pool (those not used for static translations, or for its own WAN IP address) for connections to the Internet or another outside network. When an address is no longer in use, it is returned to the address pool to be dynamically assigned later to another device.

If the addressing requirements of your network exceed the available addresses in your dynamic NAT pool, you can use the Port Address Translation (PAT) feature (also called Overloading) to associate many private addresses with one or a small group of public IP address, using port addressing to make each translation unique. With PAT enabled, the router chooses a unique port number for the IP address of each outbound translation slot. This feature is useful if you cannot allocate enough unique IP addresses for your outbound connections. Note that Port Address Translation does not occur until the address pool is depleted.

Note By default, Security Manager does not perform NAT on traffic that is meant to be transmitted over a VPN. Otherwise, any traffic appearing in both the NAT ACL and the crypto ACL defined on an interface would be sent out unencrypted because NAT is always performed before encryption. However, you can change this default setting.

Tip You can perform PAT on split-tunneled traffic on the spokes of your VPN topology directly from the Global VPN Settings page. There is no need to create a dynamic NAT rule for each spoke. Any NAT rules that you define on an individual device override the VPN setting. For more information, see Configuring VPN Global NAT Settings, page 26-42.

The Add Dynamic NAT Rule and Edit Static NAT Rule dialog boxes are used to add and edit these rules. Refer to NAT Dynamic Rule Dialog Box, page 24-11 for descriptions of the fields displayed in the table on this page.

Before You Begin

- Define the inside and outside interfaces used for NAT. See NAT Page: Interface Specification, page 24-6.

Navigation Path

- (Device view) Select **NAT** from the Policy selector, then click the **Dynamic Rules** tab.
- (Policy view) Select **NAT (Router) > Translation Rules** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Dynamic Rules** tab.

Related Topics

- NAT Policies on Cisco IOS Routers, page 24-5
- NAT Page: Static Rules, page 24-6
- NAT Page: Timeouts, page 24-13
- Standard Security Manager rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

NAT Dynamic Rule Dialog Box

Use the Add/Edit NAT Dynamic Rule dialog boxes to add or edit dynamic address translation rules. Except for their titles, the two dialog boxes are identical.

Navigation Path

Go to the NAT Page: Dynamic Rules, page 24-10 tab; click the **Add** button beneath the table to add a new rule, or select a rule in the table and click **Edit** to update that rule.

Related Topics

- Creating Access Control List Objects, page 6-53
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 24-3 NAT Dynamic Rule Dialog Box

Element	Description
Traffic Flow	In the **Access List** field, enter or Select the name of the access control list (ACL) object whose entries define the addresses requiring dynamic translation. **Note** Make sure that the specified ACL does not permit the translation of Security Manager management traffic over any device address on this router. Translating this traffic will cause a loss of communication between the router and Security Manager.
Translated Address	Use the options in this section of the dialog box to specify the method and address(es) used for dynamic translation: • **Use Interface IP** – Select this option to specify that the globally registered IP address assigned to a particular interface be used as the translated address; port addressing ensures each translation is unique. (The Enable Port Translation (Overload) option is checked automatically when you select Use Interface IP.) Enter or Select the name of the desired **Interface**. This is typically the interface from which translated packets leave the router, meaning the interface or interface role must represent an outside interface on the router (see NAT Page: Interface Specification, page 24-6). • **Address Pool** – Select this option to base address translation on the addresses you specify in the **Network Ranges** pool. Enter one or more address ranges, including the prefix, using the format `min1-max1/prefix` (in CIDR notation), where "prefix" represents a valid netmask. For example, `172.16.0.0-172.31.0.223/12`. You can add as many address ranges to the address pool as required, but all ranges must share the same prefix. Separate multiple entries with commas.

Table 24-3 NAT Dynamic Rule Dialog Box (continued)

Element	Description
Settings	This section contains two options • **Enable Port Translation (Overload)** – When selected, the router uses port addressing (PAT) if supply of global addresses in the address pool is depleted; when deselected, PAT is not used. **Note** When you use select Use Interface IP in the Translated Address section, this box is checked automatically; it cannot be changed. • **Do Not Translate VPN Traffic (Site-to-Site VPN only)** – Deselect this option to allow address translation on traffic intended for a site-to-site VPN. When selected, address translation is not performed on VPN traffic. When deselected, the router performs address translation on VPN traffic in cases of overlapping addresses between the NAT ACL and the crypto ACL. **Note** We strongly recommend that you not deselect this option, or any traffic defined in both the NAT ACL and the crypto ACL will be sent unencrypted. When you perform NAT into IPsec, we also recommend that you leave this option selected; it does not interfere with the translation of addresses arriving from overlapping networks. This setting applies only in situations where the NAT ACL overlaps the crypto ACL used by the site-to-site VPN. Because the interface performs NAT first, any traffic arriving from an address within this overlap would get translated, causing the traffic to be sent unencrypted. Leaving this box checked prevents that from happening. **Note** This option does not apply to remote access VPNs.

NAT Page: Timeouts

Use the NAT Timeouts tab of the router's NAT page to manage the timeout values for port address (overload) translations. These timeouts cause a dynamic translation to expire after a specified period of inactivity. In addition, you can use options on this page to place a limit on the number of entries allowed in the dynamic NAT table, and to modify the default timeout on all dynamic translations that do not include PAT processing.

About Dynamic NAT Timeouts

Dynamic NAT translations have a timeout period for non-use, after which they expire and are purged from the translation table. If you enable the Overload feature for performing PAT, you can specify a variety of values that provide finer control over these timeouts, because each translation entry contains additional contextual information about the traffic using it.

For example, non-DNS translations time out by default after five minutes, but DNS translations time out after 1 minute. Further, TCP translations time out after 24 hours, unless an RST or FIN is seen on the stream, in which case they time out after one minute. You can change any of these timeout values.

Note If you disable the Port Translation (Overload) feature for all dynamic rules, you need not enter any PAT-related timeout values. However, you can still modify the default timeout value for non-PAT dynamic translations. (By default, all dynamic translations expire after 24 hours.) For more information about the Overload feature, see NAT Dynamic Rule Dialog Box, page 24-11.

Navigation Path

- (Device view) Select **NAT** from the Policy selector, then click the **Timeouts** tab.
- (Policy view) Select **NAT (Router) > Translation Rules** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Timeouts** tab.

Related Topics

- NAT Page: Interface Specification, page 24-6
- NAT Page: Static Rules, page 24-6
- NAT Page: Dynamic Rules, page 24-10

Field Reference

Table 24-4 NAT Timeouts Tab

Element	Description
Max Entries	The maximum number of entries allowed in the dynamic NAT table. You can enter a value between 1 and 2147483647, or you can leave the field blank (the default), which means that the number of entries in the table is unlimited.
Timeout (sec.)	The number of seconds after which dynamic translations expire; this does not apply to PAT (overload) translations. The default is 86400 seconds (24 hours).
UDP Timeout (sec.)	The timeout value applied to User Datagram Protocol (UDP) ports. The default is 300 seconds (5 minutes). **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.
DNS Timeout (sec.)	The timeout value applied to Domain Naming System (DNS) server connections. The default is 60 seconds. **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.
TCP Timeout (sec.)	The timeout value applied to Transmission Control Protocol (TCP) ports. The default is 86400 seconds (24 hours). **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.

Table 24-4 NAT Timeouts Tab (continued)

Element	Description
FINRST Timeout (sec.)	The timeout value applied when a Finish (FIN) packet or Reset (RST) packet (both of which terminate connections) is found in the TCP stream. The default is 60 seconds. **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.
ICMP Timeout (sec.)	The timeout value applied to Internet Control Message Protocol (ICMP) flows. The default is 60 seconds. **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.
PPTP Timeout (sec.)	The timeout value applied to NAT Point-to-Point Tunneling Protocol (PPTP) flows. The default is 86400 seconds (24 hours). **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.
SYN Timeout (sec.)	The timeout value applied to TCP flows after a synchronous transmission (SYN) message (used for precise clocking) is encountered. The default is 60 seconds. **Note** This value applies only when Port Translation (Overload) is enabled for a dynamic NAT rule; see NAT Dynamic Rule Dialog Box, page 24-11.

NAT Policies on Security Devices

The following topics describe configuring network address translation (NAT) options on managed security appliances: PIX firewalls, Firewall Service Modules (FWSMs) on Catalyst switches, pre-version-8.3 Adaptive Security Appliances (ASAs), and ASA 8.3+ devices. The topics are arranged as follows:

- NAT in Transparent Mode, page 24-16
- Translation Options Page, page 24-17
- **PIX, FWSM, and pre-8.3 ASA**
  - Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
  - Address Pools, page 24-18
  - Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- **ASA 8.3+**
  - Configuring NAT on ASA 8.3+ Devices, page 24-33
  - Translation Rules: ASA 8.3+, page 24-34

NAT in Transparent Mode

Using NAT on a security appliance operating in transparent mode eliminates the need for upstream or downstream routers to perform NAT for their networks. NAT in transparent mode has the following requirements and limitations:

- When the mapped addresses are not on the same network as the transparent firewall, you need to add a static route for the mapped addresses on the upstream router that points to the downstream router (through the security appliance).
- If the real destination address is not directly connected to the security appliance, you also need to add a static route on the security appliance for the real destination address that points to the downstream router. Without NAT, traffic from the upstream router to the downstream router does not need any routes on the security appliance because it uses the MAC address table. Using NAT, however, causes the security appliance to use a route look-up instead of a MAC address look-up, so it needs a static route to the downstream router.
- Because the transparent firewall does not have any interface IP addresses, you cannot use interface PAT.
- ARP inspection is not supported. Moreover, if for some reason a host on one side of the security appliance sends an ARP request to a host on the other side of the security appliance, and the initiating host real address is mapped to a different address on the same subnet, then the real address remains visible in the ARP request.

Global Options Page

Security Manager version 4.9 supports Carrier Grade NAT to configure the block size and maximum blocks per host limit for port block allocation, for ASA devices 9.5(1) or higher. Use the Global Options page to configure these options.

Navigation Path

- (Device view) Select **NAT > Global Options** from the Device Policy selector.
- (Policy view) Select **NAT (PIX/ASA/FWSM) > Global Options** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Global Options** to create a new policy.

Related Topics

- NAT Policies on Security Devices, page 24-15
- Add and Edit NAT Rule Dialog Boxes, page 24-36

Field Reference

Table 24-5 Global Options Page

Element	Description
xlate block-allocation size	Enter a value between 32 and 4096. The default value is 512.
xlate block-allocation maximum-per-host	Enter a value between 1 to 8. The default value is 4.

Translation Options Page

Use the Translation Options page to set options that affect network address translation for the selected security appliance. These settings apply to all interfaces on the device.

Navigation Path

- (Device view) Select **NAT > Translation Options** from the Device Policy selector.
- (Policy view) Select **NAT (PIX/ASA/FWSM) > Translation Options** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Translation Options** to create a new policy.

Related Topics

- NAT Policies on Security Devices, page 24-15

Field Reference

Table 24-6 Translation Options Page

Element	Description
Enable traffic through the firewall without address translation	When selected, lets traffic pass through the security appliance without address translation. If this option is not selected, any traffic that does not match a translation rule will be dropped. **Note** This option is available only on PIX 7.x, FWSM 3.x, and ASA devices.
Enable xlate bypass	When selected, establishment of NAT sessions for untranslated traffic is disabled (this feature is called “xlate bypass”). **Note** This option is available only on FWSM 3.2 and higher. By default, the FWSM creates NAT sessions for all connections even if NAT is not used. For example, a session is created for each untranslated connection even if NAT control is not enabled, if NAT exemption or identity NAT is used, or if you use same-security interfaces and do not configure NAT. Because there is a maximum number of NAT sessions (266,144 concurrent), these kinds of NAT sessions might cause you to run into the limit. To avoid reaching the limit, enable xlate bypass. If you disable NAT control and have untranslated traffic or use NAT exemption, or if you enable NAT control and use NAT exemption, then with xlate bypass, the FWSM does not create a session for those types of untranslated traffic. However, NAT sessions are still created in the following instances: • You configure identity NAT (with or without NAT control)—identity NAT is considered to be a translation. • You use same-security interfaces with NAT control. Traffic between same-security interfaces create NAT sessions even when you do not configure NAT for the traffic. To avoid NAT sessions in this case, disable NAT control, or use NAT exemption as well as xlate bypass.

Table 24-6 Translation Options Page (continued)

Element	Description
Do not translate VPN traffic	When selected, VPN traffic passes through the security appliance without address translation.
Clear translates for existing connections	When selected, the translation slots assigned to dynamic translations and any associated connections are cleared following each session. Each session connecting through the security appliance, and undergoing some form of NAT or PAT, is assigned a translation slot known as an "xlate." These translation slots can persist after the session is complete, which can lead to a depletion of translation slots, unexpected traffic behavior, or both.

Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functions, it does not support any enhancements.

The following sections describe configuring network address translation on PIX and FWSM devices, and on pre-8.3-version ASAs. (See Configuring NAT on ASA 8.3+ Devices, page 24-33 for information about configuring NAT on ASA 8.3+ devices.)

- Address Pools, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
  - Translation Exemptions (NAT 0 ACL), page 24-20
  - Dynamic Rules Tab, page 24-22
  - Policy Dynamic Rules Tab, page 24-24
  - Static Rules Tab, page 24-26
  - General Tab, page 24-31

Address Pools

Use the Address Pools page to view and manage the global address pools used in dynamic NAT rules.

The Address Pool dialog box is used to add and edit these address pools. Refer to Address Pool Dialog Box, page 24-19 for descriptions of the fields displayed in the Global Address Pools table on this page.

Navigation Path

- (Device view) Select **NAT > Address Pools** from the Device Policy selector.
- (Policy view) Select **NAT (PIX/ASA/FWSM) > Address Pools** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Address Pools** to create a new policy.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18

Address Pool Dialog Box

Use the Address Pool dialog box to add or edit a global address pool for use in dynamic NAT rules.

Navigation Path

You open the Address Pool dialog box by clicking the Add Row or Edit Row buttons on the Address Pools, page 24-18.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18

Field Reference

Table 24-7 Address Pools Dialog Box

Element	Description
Interface Name	Enter or Select the name of the device interface on which the mapped IP addresses will be used.
Pool ID	Enter a unique identification number for this address pool, an integer between 1 and 2147483647. When configuring a dynamic NAT rule, you select a Pool ID to specify the pool of addresses to be used for translation.
IP address ranges	Enter or Select the addresses to be assigned to this address pool. You can specify these addresses as follows: • Address range for dynamic NAT (e.g., 192.168.1.1-192.168.1.15) • Subnetwork (e.g., 192.168.1.0/24) • List of addresses separated by commas (e.g., 192.168.1.1, 192.168.1.2, 192.168.1.3) • Single address to use for PAT (e.g., 192.168.1.1) • Combinations of the above (e.g., 192.168.1.1-192.168.1.15, 192.168.1.25) • Names of hosts on the connected network; these will be resolved to IP addresses.
Description	Enter a description for the address pool.
Enable Interface PAT	When checked, port address translation is enabled on the specified interface.

Translation Rules: PIX, FWSM, and pre-8.3 ASA

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any enhancements.

Use the Translation Rules page to define network address translation (NAT) rules on the selected device. The Translation Rules page consists of the following tabs:

- Translation Exemptions (NAT 0 ACL), page 24-20 – Use this tab to configure rules specifying traffic that is exempt from address translation.

Note Translation exemptions are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

- Dynamic Rules Tab, page 24-22 – Use this tab to configure dynamic NAT and PAT rules.

Note Dynamic translation rules are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

- Policy Dynamic Rules Tab, page 24-24 – Use this tab to configure dynamic translation rules based on source and destination addresses and services.

Note Policy dynamic rules are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

- Static Rules Tab, page 24-26 – Use this tab to configure static translation rules for a security appliance or shared policy.
- General Tab, page 24-31 – Use this tab to view all current translation rules, listed in the order that they will be evaluated on the device.

Note The General tab is visible only for PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules and do not need to display summary information.

Navigation Path

To access the Translation Rules page, do one of the following:

- (Device view) Select **NAT > Translation Rules** from the Device Policy selector.
- (Policy view) Select **NAT (PIX/ASA/FWSM) > Translation Rules** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Translation Rules** to create a new policy.

Translation Exemptions (NAT 0 ACL)

Use the Translation Exemptions (NAT 0 ACL) tab of the Translation Rules page to view and specify rules that exempt traffic from address translation. Rules are evaluated sequentially in the order listed. The row number indicates the rule's position in the ordering of the list. You can use the Up Row and Down Row buttons to change the position of the selected rule.

The Add/Edit Translation Exemption (NAT-0 ACL) Rule dialog box is used to add and edit these rules. Refer to Add/Edit Translation Exemption (NAT-0 ACL) Rule Dialog Box, page 24-21 for descriptions of the fields displayed in the table on this page.

Note Translation exemptions are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

Navigation Path

You can access the Translation Exemptions (NAT 0 ACL) tab from the Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19 page.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Advanced NAT Options Dialog Box, page 24-29
- General Tab, page 24-31
- Standard Security Manager rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Add/Edit Translation Exemption (NAT-0 ACL) Rule Dialog Box

Use the Add/Edit Translation Exemption (NAT-0 ACL) Rule dialog box to define and edit translation exemption rules on PIX, FWSM and pre-8.3 ASA devices in router mode, and FWSM 3.2 devices in transparent mode.

Navigation Path

You can access the Add/Edit Translation Exemption (NAT-0 ACL) Rule dialog box from the Translation Exemptions (NAT 0 ACL) tab. See Translation Exemptions (NAT 0 ACL), page 24-20 for more information.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- Advanced NAT Options Dialog Box, page 24-29

Field Reference

Table 24-8 Add/Edit Translation Exemption (NAT-0 ACL) Rule Dialog Box

Element	Description
Enable Rule	If checked, the rule is enabled. Deselect this option to disable the rule without deleting it.
Action	Select the action for this rule: • exempt – The rule identifies traffic that is exempt from NAT. • do not exempt – The rule identifies traffic that is not exempt from NAT.
Original: Interface	Enter the name of (or Select) the device interface to which the rule applies.

Table 24-8 Add/Edit Translation Exemption (NAT-0 ACL) Rule Dialog Box (continued)

Element	Description
Original: Sources	Enter IP addresses for (or Select) the source hosts and network objects to which the rule applies. Multiple entries must be separated by commas. Note that this parameter is displayed in the Translation Exemptions (NAT 0 ACL) table under the column heading "Original Address."
Translated: Direction	The rule can be applied to Inbound or Outbound traffic, as specified with this option.
Traffic flow: Destinations	Enter IP addresses for (or Select) the destination hosts and network objects to which the rule applies. Multiple entries must be separated by commas.
Category	To assign the rule to a category, choose the category from this list. Categories can help identify rules and objects using labels and color-coding. See Using Category Objects, page 6-13 for more information. **Note** No commands are generated for the Category attribute.
Description	Enter a description of the rule.
Advanced button (FWSM only)	Click to open the Advanced NAT Options Dialog Box, page 24-29 to configure advanced settings for this rule.

Dynamic Rules Tab

Use the Dynamic Rules tab of the Translation Rules page to view and configure dynamic NAT and PAT rules. Rules are evaluated sequentially in the order listed. The row number indicates the rule's position in the ordering of the list. You can use the Up Row and Down Row buttons to change the position of the selected rule.

With dynamic NAT, internal IP addresses are dynamically translated using IP addresses from a pool of global addresses. With dynamic PAT, internal IP addresses are translated to a single mapped address by using dynamically assigned port numbers with the mapped address. Dynamic translations are often used to map local RFC 1918 IP addresses to addresses that are Internet-routable.

The Add/Edit Dynamic Translation Rule dialog box is used to add and edit these rules. Refer to Add/Edit Dynamic Translation Rule Dialog Box, page 24-23 for descriptions of the fields displayed in the table on this page.

Note Dynamic translation rules are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

Navigation Path

You can access the Dynamic Rules tab from the Translation Rules page. For more information about the Translation Rules page, see Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19.

Note By default, only standard Dynamic Rule elements are displayed in this table. Additional columns for elements defined in the Advanced NAT Options dialog box can be displayed by right-clicking any column heading. (All columns are displayed by default on the General Tab, page 24-31.)

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Advanced NAT Options Dialog Box, page 24-29
- Select Address Pool Dialog Box, page 24-24
- General Tab, page 24-31
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Add/Edit Dynamic Translation Rule Dialog Box

Use the Add/Edit Dynamic Translation Rule dialog box to define and edit dynamic NAT and PAT rules.

Navigation Path

You can access the Add/Edit Dynamic Translation Rule dialog box from the Dynamic Rules tab. See Dynamic Rules Tab, page 24-22 for more information.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- Advanced NAT Options Dialog Box, page 24-29
- Select Address Pool Dialog Box, page 24-24

Field Reference

Table 24-9 Add/Edit Dynamic Translation Rule Dialog Box

Element	Description
Enable Rule	If checked, the rule is enabled. Deselect this option to disable the rule without deleting it.
Original: Interface	Enter the name or Select the device interface to which the rule applies.
Original: Address	Enter IP addresses for (or Select) the source hosts and network objects to which the rule applies. Multiple entries must be separated by commas.
Translated: Pool	Enter (or Select) the ID number of the pool of addresses used for translation; clicking Select opens the Select Address Pool Dialog Box, page 24-24. Enter a value of zero to specify this as an identity NAT rule.

Table 24-9 Add/Edit Dynamic Translation Rule Dialog Box (continued)

Element	Description
Translated: Direction	The rule can be applied to Inbound or Outbound traffic, as specified with this option.
Advanced button	Click to open the Advanced NAT Options Dialog Box, page 24-29 to configure advanced settings for this rule.

Select Address Pool Dialog Box

The Select Address Pool dialog box presents a list of global address pools; these pools are defined and managed via the Address Pools, page 24-18. Use this dialog box to select an address pool for use by a dynamic translation rule, or a policy dynamic translation rule.

Navigation Path

You can access the Select Address Pool dialog box from the Add/Edit Dynamic Translation Rule Dialog Box, page 24-23 when adding or editing a dynamic translation rule, or from the Add/Edit Policy Dynamic Rules Dialog Box, page 24-25 when adding or editing a policy dynamic translation rule.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- Address Pools, page 24-18

Field Reference

Table 24-10 Select Address Pool Dialog Box

Element	Description
Pool ID	The identification number of the address pool.
Interface	The name of the device interface to which the address pool applies.
IP Address Ranges	The IP addresses assigned to the pool; “interface” in this list indicates PAT is enabled on the specified Interface.
Description	The description provided for the address pool.
Selected Row	This field identifies the pool currently selected in the list. When you click OK to close the dialog box, this pool is assigned to the translation rule.

Policy Dynamic Rules Tab

Use the Policy Dynamic Rules tab of the Translation Rules page to view and configure dynamic translation rules based on source and destination addresses and services. Rules are evaluated sequentially in the order listed. The row number indicates the rule’s position in the ordering of the list. You can use the Up Row and Down Row buttons to change the position of the selected rule.

The Add/Edit Policy Dynamic Rule dialog box is used to add and edit these rules. Refer to Add/Edit Policy Dynamic Rules Dialog Box, page 24-25 for a description of the fields displayed in the table on this page.

Note Policy dynamic rules are only supported by PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules.

Navigation Path

You can access the Policy Dynamic Rules tab from the Translation Rules page. See Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19 for more information.

Note By default, only standard Policy Dynamic Rule elements are displayed in this table. Additional columns for elements defined in the Advanced NAT Options dialog box can be displayed by right-clicking any column heading. (All columns are displayed by default on the General Tab, page 24-31.)

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Add/Edit Policy Dynamic Rules Dialog Box, page 24-25
- Advanced NAT Options Dialog Box, page 24-29
- Select Address Pool Dialog Box, page 24-24
- General Tab, page 24-31
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Add/Edit Policy Dynamic Rules Dialog Box

Use the Add/Edit Policy Dynamic Rules dialog box to define and edit dynamic translation rules based on source and destination addresses and services.

Navigation Path

You can access the Add/Edit Policy Dynamic Rules dialog box from the Policy Dynamic Rules tab. See Policy Dynamic Rules Tab, page 24-24 for more information.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- Policy Dynamic Rules Tab, page 24-24
- Advanced NAT Options Dialog Box, page 24-29
- Select Address Pool Dialog Box, page 24-24

Field Reference

Table 24-11 Add/Edit Policy Dynamic Rules Dialog Box

Element	Description
Enable Rule	If checked, the rule is enabled. Deselect this option to disable the rule without deleting it.
Original: Interface	Enter the name of (or Select) the device interface to which the rule applies.
Original: Sources	Enter IP addresses for (or Select) the source hosts and network objects to which the rule applies. Multiple entries must be separated by commas. Note that this parameter is displayed in the Policy Dynamic Rules table under the column heading "Original Address."
Translated: Pool	Enter (or Select) the ID number of the pool of addresses used for translation; clicking Select opens the Select Address Pool Dialog Box, page 24-24. Enter a value of zero to specify this as an identity NAT rule.
Translated: Direction	The rule can be applied to Inbound or Outbound traffic, as specified with this option.
Traffic flow: Destinations	Enter IP addresses for (or Select) the destination hosts and network objects to which the rule applies. Multiple entries must be separated by commas.
Traffic flow: Services	Enter (or Select) the services to which the rule applies. Multiple entries must be separated by commas.
Category	To assign the rule to a category, choose the category from this list. Categories can help identify rules and objects using labels and color-coding. See Using Category Objects, page 6-13 for more information. **Note** No commands are generated for the Category attribute.
Description	Enter a description of the rule.
Advanced button	Click to open the Advanced NAT Options Dialog Box, page 24-29 to configure advanced settings for this rule.

Static Rules Tab

Use the Static Rules tab of the Translation Rules page to view and configure static translation rules for a security appliance or shared policy. Rules are evaluated sequentially in the order listed. The row number indicates the rule's position in the ordering of the list. You can use the Up Row and Down Row buttons to change the position of the selected rule.

With static translation, internal IP addresses are permanently mapped to a global IP address. These rules map a host address on a lower security-level interface to a global address on a higher security-level interface. For example, a static rule would be used for mapping the local address of a web server on a perimeter network to a global address that hosts on the outside interface would use to access the web server.

Caution The order of Static NAT rules on a security device is important, and Security Manager preserves this ordering during deployment. However, security appliances do not support in-line editing of Static NAT rules. This means that if you move, edit, or insert a rule anywhere above the end of the list, Security Manager will remove from the device all Static NAT rules that follow the new or modified rule, and then re-send the updated list from that point. Depending on the length of the list, this can require substantial overhead, and may result in traffic interruption. Whenever possible, add any new Static NAT rules to the end of the list.

The Add/Edit Static Rule dialog box is used to add and edit these rules. Refer to Add/Edit Static Rule Dialog Box, page 24-27 for descriptions of the fields displayed in the table on this page.

The "Nailed" Column in the Static Rules Table

In addition to the columns representing parameters specified in the Add/Edit Static Rule Dialog Box, page 24-27, the Static Rules table displays a column labeled "Nailed." This value is a product of device discovery; it cannot be changed in Security Manager.

The entry in the "Nailed" Column indicates whether TCP state tracking and sequence checking is skipped for the connection: true or false.

Navigation Path

You can access the Static Rules tab from the Translation Rules page. See Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19 for more information.

Note By default, only standard Static Rules elements are displayed in this table. Additional columns for elements defined in the Advanced NAT Options dialog box can be displayed by right-clicking any column heading. (All columns are displayed by default on the General Tab, page 24-31.)

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Add/Edit Static Rule Dialog Box, page 24-27
- Advanced NAT Options Dialog Box, page 24-29
- General Tab, page 24-31
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Add/Edit Static Rule Dialog Box

Use the Add/Edit Static Rule dialog box to add or edit static translation rules for a firewall device or shared policy.

Navigation Path

You can access the Add/Edit Static Rule dialog box from the Static Rules Tab, page 24-26.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19
- Advanced NAT Options Dialog Box, page 24-29

Field Reference

Table 24-12 Add/Edit Static Rule Dialog Box

Element	Description
Enable Rule	If checked, the rule is enabled. Deselect this option to disable the rule without deleting it.
Translation Type	Select the type of translation for this rule: NAT or PAT.
Original Interface	Enter (or Select) the device interface connected to the host or network with original addresses to be translated.
Original Address	Enter (or Select) the source address to be translated.
Translated Interface	Enter (or Select) the interface on which the translated addresses are to be used. To specify this as an identity NAT rule, enter the same interface in both this and the Original Interface fields.
Use Interface IP/Use Selected Address	Specify the address used for the Translated Interface: select Use Interface IP (address), or select Use Selected Address and enter an address, or Select a network/host object.
Enable Policy NAT	Select this option to enable Policy NAT for this translation rule.
Dest Address	If Policy NAT is enabled, specify the destination addresses of the hosts or networks to which the rule applies.
Services	If Policy NAT is enabled, enter or Select the Services to which the rule applies. **Note** For Static Policy NAT, IP is the only Service that can be specified. The syntax for service and service-object specification is: {**tcp** | **udp** | **tcp&udp**}/{*source_port_number* | *port_list_object*}/{*destination_port_number* | *port_list_object*} Note that if you enter only one port parameter, it is interpreted as the destination port (with a source port of "any"). For example, **tcp/4443** means tcp, source port any, destination port 4443, while **tcp/4443/Default Range** means tcp, source port 4443, and destination port Default Range (generally 1-65535). As with all text-entry fields, Security Manager may display auto-complete options. For example, if you type **tcp/** in this field, an auto-complete list of all Port Lists objects defined in Security Manager is displayed. This list will include system-generated objects such as DEFAULT RANGE, HTTPS and WEBPORTS. Refer to Configuring Port List Objects, page 6-102 for more information about Port Lists, and Configuring Service Objects, page 6-103 for more information about defining Services.

Table 24-12 Add/Edit Static Rule Dialog Box (continued)

Element	Description
Protocol	If PAT is the selected Translation Type, select the protocol, TCP or UDP, to which the rule applies.
Original Port	If PAT is the selected Translation Type, enter the port number to be translated. Note that this parameter is displayed in the Static Rules table under the column heading "Local Port."
Translated Port	If PAT is the selected Translation Type, enter the port number to which the original port number will be translated. Note that this parameter is displayed in the Static Rules table under the column heading "Global Port."
Category	To assign the rule to a category, choose the category from this list. Categories can help identify rules and objects using labels and color-coding. See Using Category Objects, page 6-13 for more information. **Note** No commands are generated for the Category attribute.
Description	Enter a description of the rule.
Advanced button	Click to open the Advanced NAT Options Dialog Box, page 24-29 to configure advanced settings for this rule.

Edit Translated Address Dialog Box

Use the Edit Translated Address dialog box to change just the translated address assigned to a static translation rule. The translated address is the address to which the original address is changed. The interface's IP address can be used, or you can enter a specific IP address. See Static Rules Tab, page 24-26 for more information about static rules and translated addresses.

For detailed information on editing firewall rules cells, see Editing Rules, page 12-10.

Navigation Path

Right-click the Translated Address cell in the Static Rules table (on the NAT > Translation Rules page) and choose **Edit Translated Address**.

Advanced NAT Options Dialog Box

Use the Advanced NAT Options dialog box to configure the advanced connection settings—DNS Rewrite, Maximum TCP and Maximum UDP Connections, Embryonic Limit, Timeout (PIX 6.x), and Randomize Sequence Number—for NAT and Policy NAT. You can also configure these options for Translation Exemption (NAT 0 ACL) rules on an FWSM.

Navigation Path

You can access the Advanced NAT Options dialog box by clicking the **Advanced** button when adding or editing a translation rule. See the following topics for more information:

- Add/Edit Translation Exemption (NAT-0 ACL) Rule Dialog Box, page 24-21
- Add/Edit Dynamic Translation Rule Dialog Box, page 24-23
- Add/Edit Policy Dynamic Rules Dialog Box, page 24-25
- Add/Edit Static Rule Dialog Box, page 24-27

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19

Field Reference

Table 24-13 Advanced NAT Options Dialog Box

Element	Description
Translate the DNS replies that match the translation rule	If checked, the security appliance rewrites DNS replies so an outside client can resolve the name of an inside host using an inside DNS server, and vice versa. For instance, if your NAT rule includes the real address of a host with an entry in a DNS server, and the DNS server is on a different interface from a client, then the client and the DNS server need different addresses for the host: one needs the mapped address and one needs the real address. This option rewrites the address in the DNS reply to the client. As an example, assume an inside web server, www.example.com, has the IP address 192.168.1.1, which is translated to 10.1.1.1 on the outside interface of the appliance. An outside client sends a DNS request to an inside DNS server, which will resolve www.example.com to 192.168.1.1. When the reply comes to the security appliance with DNS Rewrite enabled, the security appliance will translate the IP address in the payload to 10.1.1.1, so that the outside client will get the correct IP address. Note that the mapped host needs to be on the same interface as either the client or the DNS server. Typically, hosts that need to allow access from other interfaces use a static translation, so this option is more likely to be used with a static rule.
Max TCP Connections per Rule	Enter the maximum number of TCP connections allowed; valid values are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
Max UDP Connections per Rule	Enter the maximum number of UDP connections allowed; valid values are 0 through 65,535. If this value is set to zero, the number of connections is unlimited.
Max Embryonic Connections	Enter the number of embryonic connections allowed to form before the security appliance begins to deny these connections. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. Set this limit to prevent attack by a flood of embryonic connections. Valid values are 0 through 65,535. If this value is set to zero, the number of connections is unlimited. Any positive value enables the TCP Intercept feature. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP Intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help to minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server.

Table 24-13 Advanced NAT Options Dialog Box (continued)

Element	Description
Timeout	For PIX 6.x devices, enter a timeout value for this translation rule, in the format *hh:mm:ss*. This value overrides the default translation timeout specified in Platform > Security > Timeouts, unless this value is 00:00:00, in which case translations matching this rule use the default translation timeout (specified in Platform > Security > Timeouts).
Randomize Sequence Number	If checked, the security appliance randomizes the sequence numbers of TCP packets. Each TCP connection has two Initial Sequence Numbers (ISNs): one generated by the client and one generated by the server. The security appliance randomizes the ISN of the TCP SYN in both the inbound and outbound directions. Randomizing the ISN of the protected host prevents an attacker from predicting the next ISN for a new connection and potentially hijacking the new session. Disable this feature only if: • Another in-line security appliance is also randomizing initial sequence numbers and data is being scrambled. • You are using eBGP multi-hop through the security appliance, and the eBGP peers are using MD5. Randomization breaks the MD5 checksum. • You are using a WAAS device which requires that the security appliance not randomize the sequence numbers of connections. Disabling this option opens a security hole in the security appliance.

General Tab

Use the General tab of the Translation Rules page to view a summary of all translation rules defined for the current device or shared policy. The translation rules are listed in the order that they will be evaluated on the device.

Note The General tab is only visible for PIX, ASA and FWSM devices in router mode, and FWSM 3.2 devices in transparent mode. Other devices in transparent mode support only static translation rules and do not need to display summary information.

Navigation Path

You can access the General tab from the Translation Rules page. See Translation Rules: PIX, FWSM, and pre-8.3 ASA, page 24-19 for more information.

Related Topics

- Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18
- Translation Exemptions (NAT 0 ACL), page 24-20
- Dynamic Rules Tab, page 24-22
- Policy Dynamic Rules Tab, page 24-24
- Static Rules Tab, page 24-26

- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Field Reference

Table 24-14 General Tab - Translation Rules Summary Table

Element	Description
Note Hatching (a series of slanted lines) across an entry in the table indicates that rule is currently disabled. (See Enable Rule in Add/Edit Dynamic Translation Rule Dialog Box, page 24-23 for information about enabling and disabling these rules.)	
No.	Rules are evaluated sequentially in the order listed. This number indicates the rule's position in the ordering of the list.
Type	The type of translation rule; for example, Static, Dynamic, Exemption, etc.
Action	Displays "exempt" if the rule is exempt from NAT.
Original Interface	The ID of the device interface to which the rule is applied.
Original Address	The object names or IP addresses of the source hosts and networks to which the rule applies.
Local Port	The port number supplied by the host or network (for static PAT).
Translated Pool	The ID number of the address pool used for translation.
Translated Interface	The interface on which the translated addresses are to be used.
Translated Address	The translated addresses.
Global Port	The port number to which the original port number will be translated (for static PAT).
Destination	The object names and IP addresses of the destination hosts or networks to which the rule applies.
Protocol	The protocol to which the rule applies.
Service	The services to which the rule applies.
Direction	The traffic direction (Inbound or Outbound) on which the rule is applied.
DNS Rewrite	Whether the DNS Rewrite option is enabled: Yes or No. This option is set in the Advanced NAT Options Dialog Box, page 24-29.
Maximum TCP Connections	The maximum number of TCP connections allowed to connect to the statically translated IP address. If zero, the number of connections is unlimited. This option is set in the Advanced NAT Options Dialog Box, page 24-29.
Embryonic Limit	The number of embryonic connections allowed to form before the security appliance begins to deny these connections. If zero, the number of connections is unlimited. A positive number enables the TCP Intercept feature. This option is set in the Advanced NAT Options Dialog Box, page 24-29.

Table 24-14 General Tab - Translation Rules Summary Table (continued)

Element	Description
Maximum UDP Connections	The maximum number of UDP connections allowed to connect to the statically translated IP address. If zero, the number of connections is unlimited. This option is set in the Advanced NAT Options Dialog Box, page 24-29.
Timeout	For PIX 6.x devices, this is the timeout value for a static translation rule. This value overrides the default translation timeout specified in Platform > Security > Timeouts. A Timeout value of 00:00:00 here means that translations matching this rule should use the default translation timeout specified in Platform > Security > Timeouts.
Randomize Sequence Number	Whether the security appliance will randomize the sequence number of TCP packets: Yes or No. This option is set in the Advanced NAT Options Dialog Box, page 24-29, and is enabled by default.
Category	The category to which the rule is assigned. Categories use labels and color-coding to help identify rules and objects. See Using Category Objects, page 6-13 for more information. **Note** No commands are generated for the Category attribute.
Description	The description of the rule, if provided.
Last Ticket(s)	Shows the ticket(s) associated with last modification to the rule. You can click the ticket ID in the Last Ticket(s) column to view details of the ticket and to navigate to the ticket. If linkage to an external ticket management system has been configured, you can also navigate to that system from the ticket details (see Ticket Management Page, page 11-72).

Configuring NAT on ASA 8.3+ Devices

The following section describes configuring network address translation on version 8.3 or later ASA devices:

- Translation Rules: ASA 8.3+, page 24-34
  - Add and Edit NAT Rule Dialog Boxes, page 24-36
  - Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42
- Per-Session NAT Rules: ASA 9.0(1)+

See Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18 for information about configuring NAT on other security appliances. Refer to About "Simplified" NAT on ASA 8.3+ Devices, page 24-4 for general information about NAT rules, and the changes to NAT configuration implemented on the ASA 8.3.

Note You can create a NAT object only if you have the Modify privilege mapped to your role. Cisco Security Manager displays error message for authorization.

Translation Rules: ASA 8.3+

Use the Translation Rules page to manage network address translation (NAT) rules on the selected ASA 8.3+ device. See NAT Policies on Security Devices, page 24-15 for information about configuring Translation Rules on other security devices.

Two types of NAT rules are displayed in this table: "manual" rules added by you and any other users, and "automatic" rules generated and applied by Security Manager when an object with NAT properties is assigned to the device. These are referred to as "NAT rules" and "Network Object NAT rules," respectively.

Some Features of the Translation Rules Table

This Translation Rules table is a standard Security Manager rules table, as described in Using Rules Tables, page 12-8. For example, you can move, show and hide columns; you can re-order the manual rules; and you can right-click certain table cells to edit that parameter. In addition, the following features are specific to this Translation Rules table:

- All rules are assigned to one of three pre-defined sections in the table:
  - **NAT Rules Before** – These are rules you or another user have "manually" defined on the device. You can specify that a rule be added to this section by clicking the section heading before adding the rule, although if you do not specify a section, the new rule will be added to this section by default.
  - **Network Object NAT Rules** – These are rules generated and ordered automatically by Security Manager when network objects that include NAT properties are assigned to the device. See Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42 for information about assigning NAT properties to objects. See the section "The NAT Table" in About "Simplified" NAT on ASA 8.3+ Devices, page 24-4 for information about how these rules are ordered.

    Note This section is not displayed in the Translation Rules table in Policy View because these rules are device-specific.

  - **NAT Rules After** – These also are rules you or another user have manually defined on the device. You can specify that a rule is added to this section by clicking the section heading before adding the rule.

  The NAT rules listed in this table are processed on a first-match basis; therefore, order is important. Providing a manual section both before and after the automatic rules lets you ensure all your rules are in the appropriate order, since you can re-order rules only within their section. The rules in each section take precedence over the rules in the section below it. For example, the rules in the top, "Before" section take precedence over the rules in the Network Object NAT section, and so on.
- The type of each rule—Static, Dynamic PAT, or Dynamic NAT and PAT— is indicated visually in the table by presenting ***(S)***, ***(DP)***, or ***(DNP)*** in blue following the Source parameter in the "Translated" column.
- A Bi-directional rule is a static rule that actually consists of two paired rules, one each for outgoing and incoming translation of the specified source and destination values. Each Bi-directional rule entry in the rules table is presented as two lines.

  For example, if Bi-directional is chosen when you create a static rule with Host1 in the Source field and Host2 in the Translated field, two lines are added to the rules table: one with Host1 being translated to Host2, and one with Host2 being translated to Host1.

Related Topics

- NAT Policies on Security Devices, page 24-15
- About "Simplified" NAT on ASA 8.3+ Devices, page 24-4
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Navigation Path

- (Device view) Select **NAT > Translation Rules** from the Device Policy selector.
- (Policy view) Select **NAT (PIX/ASA/FWSM) > Translation Rules** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Translation Rules** to create a new policy.

The Translation Rules page is displayed. Note that in Policy View, the Network Object NAT Rules section is not displayed because those rules are device-specific.

Adding, Editing and Deleting Rules

To **add** a NAT rule:

1. Select the heading of the section to which the rule is to be added. If you do not select a heading, the rule will be added to *NAT Rules Before* by default.
2. Open the Add NAT Rule dialog box: either click the Add Row button at the bottom of the table, or right-click anywhere in the table (except on an existing rule entry) and choose Add Row from the pop-up menu.
3. Define the rule and then click OK to close the dialog box, adding the rule to the table.

To **edit** a NAT rule:

1. Open the Edit NAT Rule dialog box for the desired rule: either select the rule in the NAT rules table and then click the Edit Row button at the bottom of the table, or simply right-click the desired rule entry and choose Edit Row from the pop-up menu.
2. Edit the rule and then click OK to close the dialog box.

See Add and Edit NAT Rule Dialog Boxes, page 24-36 for a complete description of the Add NAT Rule dialog box.

To **delete** a NAT rule, select the rule in the table and click the Delete Row button at the bottom of the table, or simply right-click the desired rule entry and choose Delete Row from the pop-up menu.

Note To remove a Network Object NAT rule from this table, you must uncheck the Add Automatic Address Translation NAT Rule option, or change the device to which the rule is assigned, in the related Edit Network Host dialog box. See Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42 for additional information.

Enabling and Disabling Rules

You can disable one or more consecutive rules without removing them from the table, as follows:

1. Select the rule(s) to be disabled. If selecting a contiguous block of rules, click the first and then Shift-click the last rule of the block.
2. Right-click a selected rule, and choose **Disable** from the pop-up menu.

Disabled rules are grayed-out in the table.

To re-enable one or one or more consecutive disabled rules, repeat this process, choosing **Enable** from the pop-up menu.

Add and Edit NAT Rule Dialog Boxes

Use the Add NAT Rule dialog box to add a NAT rule to the selected ASA 8.3+ device; this dialog box is not available on earlier-version ASAs, nor on PIX or FWSM devices. Refer to Configuring NAT on PIX, FWSM, and pre-8.3 ASA Devices, page 24-18 for information about adding and editing NAT rules on those devices.

Note Except for their titles, the Add NAT Rule and Edit NAT Rule dialog boxes are identical, and the following descriptions apply to both.

Navigation Path

To add a rule, select the section to which you want the rule added (NAT Rules Before or NAT Rules After), and then click the Add Row button below the rules table, or right-click anywhere inside the table and choose **Add Row** to open the Add NAT Rule dialog box. Note that if you do not select a section, the new rule is added to the NAT Rules Before section.

To edit a rule, select the rule and click the Edit Row button, or simply right-click the rule and choose the Edit Row command, to open the Edit NAT Rule dialog box for that rule.

Related Topics

- Chapter 24, “Configuring Network Address Translation”
- Translation Rules: ASA 8.3+, page 24-34
- Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42

Field Reference

Table 24-15 Add and Edit NAT Rule Dialog Boxes

Element	Description
Source Interface	The name of the interface on which a packet may originate; this is the “real” interface. Defaults to “any,” which represents all interfaces. Enter or Select the desired interface. **Note** In transparent firewall mode, you must set specific interfaces.
Destination Interface	**Destination Interface** – The name of the interface on which a packet may terminate; this is the “mapped” interface. Defaults to “any,” which represents all interfaces. Enter or Select the desired interface. **Note** In transparent firewall mode, you must set specific interfaces.

Table 24-15 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
Source NAT Type	The type of translation rule you are creating: • **Static** – Provides static assignment of real addresses to mapped addresses. • **Dynamic PAT (Hide)** – Provides dynamic assignment of multiple local addresses to a single global IP address and a unique port number, in effect "hiding" the local addresses behind the one global address. • **Dynamic NAT and PAT** – Provides dynamic assignment of real addresses to mapped addresses, and real ports to mapped ports. Selecting this option adds the PAT Pool Address Translation options to the dialog box. On devices operating in routed mode, this option also provides the fallthrough option described below. **Note** This selection applies only to the specified source translation; destination translation is always static.
Source Translation	
Original Source	The source address the NAT rule will translate. If this is a range or network, all addresses in the range or network are translated.
Translated Source Address Interface	Whether the translation is based on an address or an interface on the device. Select either: • **Address** – Translate the original address using the Networks/Hosts object specified in the Translated Source field. This entry represents the pool of translation addresses: enter or Select the desired Networks/Hosts; defaults to the Original Source (which will produce an Identity NAT rule). • **Interface** – Translate the original address based on the interface specified in the Destination Interface field. For port address translation based on this interface, be sure to configure the options in the Service Translation section (in the Advanced panel of this dialog box). If the Destination Interface is not defined, the Address/Interface selection reverts to Address and the Original Source is inserted into the Address field. This produces an Identity NAT rule, meaning the specified address(es) are translated to themselves (effectively not translated); Identity NAT applies to outbound connections only. **Note** These options are not available when Dynamic NAT and PAT is the selected Type, nor are they available on devices operating in transparent mode.

Table 24-15 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
PAT Pool Address Translation	This option is available when Dynamic NAT and PAT is the selected Type. The related parameters let you specify a "pool" of IP addresses to be used for specifically for port address translation, as well as change the algorithm used for PAT mapping. Refer to PAT Pools and Round Robin Allocation, page 24-41 for additional information about these features. Check the PAT Pool Address Translation box to enable the following options: • **Address** or **Interface** – Select Address to indicate that the PAT Pool Address field contains networks/hosts (or networks/hosts objects) for use as the PAT pool. Select Interface to provide a Fallthrough Interface. • **Address** – Enter or Select the desired Networks/Hosts or desired Interface according to your Address or Interface selection above. • **Use Round Robin Allocation** – Check this box to map addresses/ports using a "round-robin" approach. See PAT Pools and Round Robin Allocation, page 24-41 for more information about this option. • **Extended PAT Table** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to enable extended PAT. Extended PAT uses 65535 ports per service, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. This option is available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1). • **Flat Port Range** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to enable use of the entire 1024 to 65535 port range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is not available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also select **Include Reserve Ports**. • **Include Reserve Ports** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to include the reserve ports, 1-1023, in the PAT range. • **Block Allocation** (Available for ASA 9.5(1) and later) – Check this box to allocate a block of ports per host. This feature is supported from Security Manager version 4.9 onwards for ASA devices 9.5(1) or higher.

Table 24-15 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
Destination Translation	
Use the options in this section to configure optional static translation of destination addresses: **Note** If defined, Destination Translation is always static, regardless of the rule Type. **Note** These options are not available on devices operating in transparent mode.	
Original Destination Address Interface	Whether the translation is based on an address or an interface on the device. Select either: • **Address** - Translate the original destination using the Networks/Hosts object specified in the Translated Destination field. If Address is selected, specify the Networks/Hosts object, whose original destination addresses should be translated, in the **Original Destination** entry field. • **Interface** – Translate the original destination using the Networks/Hosts object specified in the Translated Destination field. If Interface is selected, enter or select the desired interface in the Destination Interface field. The Interface Selector list contains all interfaces currently defined on the device.
Translated Destination	This entry represents the pool of destination addresses to use for translation: enter or select the desired Networks/Hosts object.
Service Translation	
Use the options in this section to configure port address translation. These service objects represent a service protocol (TCP or UDP), and one or more ports. The mapping of original ports to translated port is circular. That is, the first original value is mapped to the first translated value, and the second original value is mapped to the second translated value, and so on until all original values are translated. If the pool of translated port is exhausted before that point, mapping continues using the first translated value again. See Understanding and Specifying Services and Service and Port List Objects, page 6-100 for information about configuring service objects. **Note** Service Translation and the following **Translate DNS replies that match this rule** option cannot be used together.	
Original Service	Enter or select the Service object that defines the service(s) to be translated. Leave the Original Service field blank to configure translation of any service to the specified Translated Service. **Note** The protocol specified in both Service objects must be the same.
Translated Service	Enter or select the Service object that provides the service(s) to be used for translation.
Options	

Table 24-15 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
Translate DNS replies that match this rule	When checked, addresses embedded in DNS replies that match this rule are rewritten. For DNS replies traversing from a mapped interface to a real interface, the Address (or "A") record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the A record is rewritten from the real value to the mapped value. Note that DNS inspection must be enabled to support this functionality.
Fallthrough to Interface PAT (Destination Interface)	When checked, dynamic PAT back-up is enabled. When the pool of dynamic NAT addresses is depleted, port address translation is performed, using the address pool specified in the Use Address field. This option is available only when Dynamic NAT and PAT is the chosen Type on devices operating in routed mode.
IPv6	When selected, the IPv6 address of the interface is used.
Net to net mapping of IPv4 to IPv6	When checked, translates the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For a one-to-one translation, you must select this option.
Do not proxy ARP on Destination Interface	Check this box to disable proxy ARP on the specified Destination Interface. This option is available only when Static is the chosen rule Type. **Note** This option is available on ASA 8.4.2+ devices, only when Bidirectional is the chosen Direction. By default, all NAT rules include proxy ARP on the egress interface. A NAT Exempt rule is used to bypass NAT for both ingress and egress traffic, relying on route look-up to locate the egress interface. Thus, Proxy ARP should be disabled for NAT Exempt rules. (The NAT Exempt rules always take priority and appear above all other NAT rules in the Translation Rules table.) **Note** You also can disable Proxy ARP on individual interfaces, as described in Configuring No Proxy ARP, page 56-1.
Perform route lookup for Destination Interface	If this option is selected, the egress interface is determined using route look-up instead of using the specified Destination Interface. Be sure this box is checked for a NAT Exempt rule. This option is supported only for Static Identity NAT. **Note** This option is available on ASA 8.4.2+ devices, only when Bidirectional is the chosen Direction. The option is not available on devices operating in transparent mode.

Table 24-15 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
Unidirectional	This feature lets you configure a static NAT rule in a single direction only; or dual rules, one each for both directions (forward and reverse). When selected, a single static NAT is created, as specified by the other options in this dialog box. Dynamic rules are uni-directional by default. If deselected, two linked static NAT rules are created, encompassing both directions of the translation, as specified by the other options in this dialog box. Note that each bi-directional rule entry in the rules table consists of two lines.
Description	(Optional) Provide a description of the rule.
Category	(Optional) Choose a category to assign to the rule. Categories can help you organize and identify rules and objects; see Using Category Objects, page 6-13 for more information. **Note** This option is not available when Dynamic NAT and PAT is the chosen rule Type.

PAT Pools and Round Robin Allocation

Adaptive Security Appliances, version 8.4.2 and later, include two features that let you alter how port address translation (PAT) occurs: you can explicitly define a pool of IP addresses specifically for PAT, and you can select a "round robin" algorithm for port allocation during PAT.

These features simplify configuration of large numbers of PAT addresses, and help prevent a large number of connections from a single PAT address, which can appear to be part of a DoS attack.

Explicit PAT Pool Definition

Prior to version 8.4.2, when you defined a Dynamic NAT and PAT rule, you provided a "pool" of IP addresses (in the Translated Source field of the Add/Edit NAT Rule dialog boxes) to be used for translation. This pool could consist of individual IP addresses, ranges of addresses, Networks/Hosts objects, or Network/Host group objects, and combinations thereof.

Ranges and objects with more than one IP address were considered to be in the "NAT Pool," while individual IP addresses and group objects consisting of one or more individual addresses were considered to be part of the "PAT Pool."

Address translation on the device would work its way through the NAT Pool until all available addresses were exhausted. Port address translation would then begin using the PAT Pool—assigning ports on the first IP address in the PAT Pool until all ports (approximately 64,000) are assigned, then assigning ports on the next address in the pool, and so on. When all ports are fully subscribed on all IP addresses in the PAT Pool, no further translation could occur.

On version 8.4.2 and later ASA devices, you can explicitly define a separate PAT Pool for a Dynamic NAT and PAT rule. If you do so, the first collection of addresses (defined in the Translated Source field) is considered the NAT Pool, while the PAT Pool addresses are specified in the PAT Pool Address Translation field.

Note If you do not explicitly specify a PAT Pool, address translation takes place as described for pre-8.4.2 devices.

Refer to Add and Edit NAT Rule Dialog Boxes, page 24-36 for more information about the defining translation rules.

Round Robin Port Assignment

On version 8.4.2 and later ASA devices, you also can specify an alternate method of port assignment during PAT processing. As mentioned earlier, PAT port numbers are assigned to a single IP address in succession until the final port number is assigned, and then the process begins again with the next available IP address in the pool.

However, a new parameter on 8.4.2 and later devices—Use Round Robin Allocation for PAT Pool—lets you specify "round robin" cycling through available IP addresses and port numbers. This method assigns an address/port combination using each successive address in the pool; it then uses the first address again with a different port, proceeds to the second address again, and so on.

Further, the round-robin algorithm incorporates two additional principles it will attempt to adhere to when assigning address/port combinations during PAT processing:

- If a specific source-to-destination mapping already exists, the algorithm attempts to use the existing translation for the new connection. If this is not possible (for example, when all ports for that IP address have been exhausted), the algorithm proceeds with standard round-robin cycling.
- If possible, the original source port number is used as the mapped port number. That is, if the port number of the address/port combination to be translated is 4904, for example, and 4904 is available with the next IP address in the PAT Pool, the translated address will be *PAT_address*/4904. Note if this is not possible (that port is not available with the next PAT address), the algorithm proceeds with standard round-robin cycling.

Note If you do not explicitly specify Round Robin Allocation, port-allocation cycling occurs as described for pre-8.4.2 devices.

Add or Edit Network/Host Dialog Box: NAT Tab

Use the NAT tab in any of the dialog boxes used to add or edit host, network, or address range objects to create or update object NAT rules. This NAT configuration is used only for ASA 8.3+ devices; if you use the object on any other type of device, the NAT configuration is ignored.

The NAT configuration is created as a device override and is not kept in the global object. Therefore, you must select the **Allow Value Override per Device** option if you configure these NAT options. (This option is selected automatically when you close the dialog box.)

This topic describes the fields on the NAT tab. For information about the fields on the General tab, see Add or Edit Network/Host Dialog Box, page 6-83.

Navigation Path

Select the NAT tab on the Add or Edit Network/Host Dialog Box when creating or editing a host, network, or address range object.

Related Topics

- Chapter 24, "Configuring Network Address Translation"
- Creating Networks/Hosts Objects, page 6-82
- Understanding Networks/Hosts Objects, page 6-80
- Specifying IP Addresses During Policy Definition, page 6-87
- Policy Object Manager, page 6-4

Field Reference

Table 24-16 Network/Host Dialog Box NAT Tab

Element	Description
Add Automatic Address Translation NAT Rule	If checked, a network address translation (NAT) rule, as defined here, will be applied to the device specified in the Translated By field. The rule will appear in the Network Object NAT Rule section of the Translation Rules table for that device (see Translation Rules: ASA 8.3+, page 24-34).
Translated By	The device on which you are configuring the NAT rule. Click Select to select the device from a list. The list is filtered to show only ASA 8.3+ devices.
Source Interface	The name of the interface on which a packet may originate; this is the "real" interface. Defaults to **any**, which represents all interfaces.
Destination Interface	The name of the interface on which a packet may terminate; this is the "mapped" interface. Defaults to **any**, which represents all interfaces.
Type	The type of translation rule you are creating: • **Static** – Enables static assignment of real addresses to mapped addresses. • **PAT (Hide)** – Enables dynamic assignment of multiple local addresses to a single global IP address and a unique port number. • **Dynamic NAT and PAT** – Enables dynamic assignment of real addresses to mapped addresses, and real ports to mapped ports.
Source Translation	
Original value	Shows the address configured on the General tab of this dialog box. This is the source address the NAT rule will translate. If it is a range or network, all addresses in the range or network will be translated.
Translated Source Use Address Use Interface (available only for Static and PAT)	Whether the translation is based on an address or an interface on the device: • **Use Address**—Translate the original address using the specified address or network/host object. Enter the address or object name in the **Address** field, or click Select to select the object from a list. • **Use Interface** – Translate the original address based on the interface specified in the Destination Interface field. **Note** The Use Interface options are available only when either Static or PAT (Hide) is chosen as the Type.

Table 24-16 Network/Host Dialog Box NAT Tab (continued)

Element	Description
PAT Pool Address Translation	This option is available when Dynamic NAT and PAT is the selected Type. The related parameters let you specify a "pool" of IP addresses to be used for specifically for port address translation, as well as change the algorithm used for PAT mapping. Refer to PAT Pools and Round Robin Allocation, page 24-41 for additional information about these features. Check the PAT Pool Address Translation box to enable the following options: • **Use Address** or **Use Interface** – Select Use Address to indicate that the PAT Pool Address field contains networks/hosts (or networks/hosts objects) for use as the PAT pool. Select Use Interface to provide a Fallthrough Interface. • **PAT Pool Address** – Enter or Select the desired Networks/Hosts or desired Interface according to your Address or Interface selection above. • **Use Round Robin Allocation for PAT Pool** – Check this box to map addresses/ports using a "round-robin" approach. See PAT Pools and Round Robin Allocation, page 24-41 for more information about this option. • **Extended PAT Table** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to enable extended PAT. Extended PAT uses 65535 ports per service, as opposed to per IP address, by including the destination address and port in the translation information. Normally, the destination port and address are not considered when creating PAT translations, so you are limited to 65535 ports per PAT address. For example, with extended PAT, you can create a translation of 10.1.1.1:1027 when going to 192.168.1.7:23 as well as a translation of 10.1.1.1:1027 when going to 192.168.1.7:80. This option is available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1). • **Flat Port Range** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to enable use of the entire 1024 to 65535 port range when allocating ports. When choosing the mapped port number for a translation, the ASA uses the real source port number if it is available. However, without this option, if the real port is not available, by default the mapped ports are chosen from the same range of ports as the real port number: 1 to 511, 512 to 1023, and 1024 to 65535. To avoid running out of ports at the low ranges, configure this setting. To use the entire range of 1 to 65535, also select **Include Reserve Ports**. • **Include Reserve Ports** (Available for ASA 8.4(3) and later, not including 8.5(1) or 8.6(1)) - Check this box to include the reserve ports, 1-1023, in the PAT range.

Table 24-16 Network/Host Dialog Box NAT Tab (continued)

Element	Description
Service Translation	
Use the options in this section of the Advanced panel to configure static port address translation: (Available for Static rules only.) **Note** Service Translation and the **Translate DNS replies that match this rule** option cannot be used together.	
Protocol	Whether a TCP or UDP port.
Original Port	The port on which the traffic enters the device.
Translated Port	The port number which is to replace the original port number.
Options	
Translate DNS replies that match this rule	When checked, addresses embedded in DNS replies that match this rule are rewritten. For DNS replies traversing from a mapped interface to a real interface, the Address (or "A") record is rewritten from the mapped value to the real value. Conversely, for DNS replies traversing from a real interface to a mapped interface, the A record is rewritten from the real value to the mapped value. Note that DNS inspection must be enabled to support this functionality. **Note** This option and Service Translation cannot be used together.
Fallthrough to Interface PAT (Destination Interface)	When checked, dynamic PAT back-up is enabled. When the pool of dynamic NAT addresses is depleted, port address translation is performed, using the address pool specified in the Use Address field. This option is available only when Dynamic NAT and PAT is the chosen Type on devices operating in routed mode.
IPv6	When selected, the IPv6 address of the interface is used.
Net to net mapping of IPv4 to IPv6	When checked, translates the first IPv4 address to the first IPv6 address, the second to the second, and so on. Without this option, the IPv4-embedded method is used where the 32-bits of the IPv4 address is embedded after the IPv6 prefix. For a one-to-one translation, you must select this option.
Do not proxy ARP on Destination Interface	Check this box to disable proxy ARP on the specified Destination Interface. This option is available only when Static is the chosen rule Type. By default, all NAT rules include proxy ARP on the egress interface. A NAT Exempt rule is used to bypass NAT for both ingress and egress traffic, relying on route look-up to locate the egress interface. Thus, Proxy ARP should be disabled for NAT Exempt rules. (The NAT Exempt rules always take priority and appear above all other NAT rules in the Translation Rules table.) **Note** You also can disable Proxy ARP on individual interfaces, as described in Configuring No Proxy ARP, page 56-1.

Table 24-16 Network/Host Dialog Box NAT Tab (continued)

Element	Description
Perform route lookup for Destination Interface	If this option is selected, the egress interface is determined using route look-up instead of using the specified Destination Interface. Be sure this box is checked for a NAT Exempt rule. This option is supported only for Static Identity NAT. **Note** This option is not available on devices operating in transparent mode.

Per-Session NAT Rules: ASA 9.0(1)+

Use the Per-Session NAT Rules page to configure per-session PAT rules on the selected ASA 9.0(1)+ device. By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. You can configure per-session rules to use multi-session PAT for specific traffic.

Per-Session PAT vs. Multi-Session PAT (Version 9.0(1) and Later)

The per-session PAT feature improves the scalability of PAT and, for clustering, allows each member unit to own PAT connections; multi-session PAT connections have to be forwarded to and owned by the control unit. At the end of a per-session PAT session, the ASA sends a reset and immediately removes the xlate. This reset causes the end node to immediately release the connection, avoiding the TIME_WAIT state. Multi-session PAT, on the other hand, uses the PAT timeout, by default 30 seconds. For "hit-and-run" traffic, such as HTTP or HTTPS, the per-session feature can dramatically increase the connection rate supported by one address. Without the per-session feature, the maximum connection rate for one address for an IP protocol is approximately 2000 per second. With the per-session feature, the connection rate for one address for an IP protocol is 65535/average-lifetime.

By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate. For traffic that can benefit from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT by creating a per-session deny rule.

Some Features of the Per-Session NAT Rules Table

This Translation Rules table is a standard Security Manager rules table, as described in Using Rules Tables, page 12-8. For example, you can move, show and hide columns; you can re-order the rules; and you can right-click certain table cells to edit that parameter.

The NAT rules listed in this table are processed on a first-match basis; therefore, order is important.

Related Topics

- Add and Edit Per Session NAT Rule Dialog Boxes, page 24-47
- NAT Policies on Security Devices, page 24-15
- About “Simplified” NAT on ASA 8.3+ Devices, page 24-4
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Navigation Path

- (Device view) Select **NAT > Per-Session NAT Rules** from the Device Policy selector.

- (Policy view) Select **NAT (PIX/ASA/FWSM) > Per-Session NAT Rules** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **Translation Rules** to create a new policy.

The Per-Session NAT Rules page is displayed.

Adding, Editing and Deleting Rules

To **add** a per-session NAT rule:

1. Select the rule under which the rule is to be added. If you do not select a heading, the rule will be added to the end of the table by default.
2. Open the Add Per-Session NAT Rule dialog box: either click the **Add Row** button at the bottom of the table, or right-click anywhere in the table and choose **Add Row** from the pop-up menu.
3. Define the rule and then click **OK** to close the dialog box, adding the rule to the table.

See Add and Edit Per Session NAT Rule Dialog Boxes, page 24-47 for a complete description of the Add Per-Session NAT Rule dialog box.

To **edit** a per-session NAT rule:

1. Open the Edit Per-Session NAT Rule dialog box for the desired rule: either select the rule in the Per-Session NAT rules table and then click the **Edit Row** button at the bottom of the table, or simply right-click the desired rule entry and choose **Edit Row** from the pop-up menu.
2. Edit the rule and then click **OK** to close the dialog box.

See Add and Edit Per Session NAT Rule Dialog Boxes, page 24-47 for a complete description of the Edit Per-Session NAT Rule dialog box.

To **delete** a per-session NAT rule, select the rule in the table and click the **Delete Row** button at the bottom of the table, or simply right-click the desired rule entry and choose **Delete Row** from the pop-up menu.

Enabling and Disabling Rules

You can disable one or more consecutive rules without removing them from the table, as follows:

1. Select the rule(s) to be disabled. If selecting a contiguous block of rules, click the first and then Shift-click the last rule of the block.
2. Right-click a selected rule, and choose **Disable** from the pop-up menu.

   Disabled rules are grayed-out in the table.

To re-enable one or one or more consecutive disabled rules, repeat this process, choosing **Enable** from the pop-up menu.

Add and Edit Per Session NAT Rule Dialog Boxes

By default, all TCP PAT traffic and all UDP DNS traffic uses per-session PAT. To use multi-session PAT for traffic, you can configure per-session PAT rules: a permit rule uses per-session PAT, and a deny rule uses multi-session PAT.

For more information about per-session vs. multi-session PAT, see Per-Session NAT Rules: ASA 9.0(1)+.

Defaults

By default, the following rules are installed:

- Permit TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6)

- Permit UDP from any (IPv4 and IPv6) to domain

These rules do not appear in the rule table.

Note You cannot remove these rules, and they always exist after any manually-created rules. Because rules are evaluated in order, you can override the default rules. For example, to completely negate these rules, you could add the following:

Deny TCP from any (IPv4 and IPv6) to any (IPv4 and IPv6)
Deny UDP from any (IPv4 and IPv6) to domain

Navigation Path

From the Per-Session NAT Rules: ASA 9.0(1)+ page, do one of the following:

- To add a rule, select the rule under which you want the rule added, and then click the **Add Row** button below the rules table, or right-click anywhere inside the table and choose **Add Row** to open the Add Per-Session NAT Rule dialog box.
- To edit a rule, select the rule and click the **Edit Row** button, or simply right-click the rule and choose **Edit Row** to open the Edit Per-Session NAT Rule dialog box for that rule.

Related Topics

- Per-Session NAT Rules
- Chapter 24, "Configuring Network Address Translation"
- Translation Rules: ASA 8.3+, page 24-34
- Add or Edit Network/Host Dialog Box: NAT Tab, page 24-42

Field Reference

Table 24-17 Add and Edit NAT Rule Dialog Boxes

Element	Description
Action	The action for this rule: Permit or Deny. A permit rule uses per-session PAT; a deny rule uses multi-session PAT.
Original Network	The source address or addresses (or Networks/Hosts objects) to which the rule applies. If this is a range or network, all addresses in the range or network are translated.
Destination Network	The destination address or addresses (or Networks/Hosts objects) to which the rule applies.
Service (tcp/udp Only)	Enter or Select the Service object that defines the service(s) to be translated. These service objects represent a service protocol (TCP or UDP), and one or more ports. See Understanding and Specifying Services and Service and Port List Objects, page 6-100 for information about configuring service objects.

Table 24-17 Add and Edit NAT Rule Dialog Boxes (continued)

Element	Description
Category	(Optional) Choose a category to assign to the rule. Categories can help you organize and identify rules and objects; see Using Category Objects, page 6-13 for more information. **Note** This option is not available when Dynamic NAT and PAT is the chosen rule Type.
Description	(Optional) Provide a description of the rule.

PART 3

VPN Configuration

Managing Site-to-Site VPNs: The Basics

A virtual private network (VPN) consists of multiple remote peers transmitting private data securely to one another over an unsecured network, such as the Internet. Site-to-site VPNs use tunnels to encapsulate data packets within normal IP packets for forwarding over IP-based networks, using encryption to ensure privacy and authentication to ensure integrity of data.

In Cisco Security Manager, site-to-site VPNs are implemented based on IPsec policies that are assigned to VPN topologies. An IPsec policy is a set of parameters that define the characteristics of the site-to-site VPN, such as the security protocols and algorithms that will be used to secure traffic in an IPsec tunnel. Security Manager translates IPsec policies into CLI commands that can be deployed to the devices in the VPN topology. Several policy types might be required to define a full configuration image that can be assigned to a VPN topology, depending on the IPsec technology type.

The Site-to-Site VPN Manager defines and configures site-to-site VPN topologies and policies on Cisco IOS security routers, PIX Firewalls, Catalyst VPN Service Modules, and Adaptive Security Appliance (ASA) firewall devices.

Tip In ASA documentation, site-to-site VPNs are called LAN-to-LAN VPNs. These phrases are equivalent, and we use “site-to-site VPN” in this documentation.

You can access the Site-to-Site VPN Manager by selecting **Manage > Site-To-Site VPNs** or clicking the Site-To-Site VPN Manager button on the toolbar.

You can also configure shared policies in Policy view and view and configure topologies in Device view. In Policy View, you can assign IPsec policies to VPN topologies.

This chapter contains the following topics:

- Understanding VPN Topologies, page 25-2
- Understanding IPsec Technologies and Policies, page 25-5
- Accessing Site-to-Site VPN Topologies and Policies, page 25-17
- Site-To-Site VPN Discovery, page 25-20
- Creating or Editing VPN Topologies, page 25-28
- Creating or Editing Extranet VPNs, page 25-66
- Deleting a VPN Topology, page 25-71

Understanding VPN Topologies

A VPN topology specifies the peers and the networks that are part of the VPN and how they connect to one another. After you create a VPN topology, the policies that can be applied to your VPN topology become available for configuration, depending on the assigned IPsec technology.

Security Manager supports three main types of topologies—hub and spoke, point to point, and full mesh, with which you can create a site-to-site VPN. Not all policies can be applied to all VPN topologies. The policies that can be applied depend on the IPsec technology that is assigned to the VPN topology. In addition, the IPsec technology that is assigned to a VPN depends on the topology type. For example, the DMVPN and Easy VPN technologies can only be applied in a hub-and-spoke topology.

For more information, see Understanding IPsec Technologies and Policies, page 25-5.

The following topics describe:

- Hub-and-Spoke VPN Topologies, page 25-2
- Point-to-Point VPN Topologies, page 25-3
- Full Mesh VPN Topologies, page 25-4
- Implicitly Supported Topologies, page 25-5

Hub-and-Spoke VPN Topologies

In a hub-and-spoke VPN topology, multiple remote devices (spokes) communicate securely with a central device (hub). A separate, secured tunnel extends between the hub and each individual spoke.

The following illustration shows a typical hub-and-spoke VPN topology.

Figure 25-1 Hub-and-Spoke VPN Topology

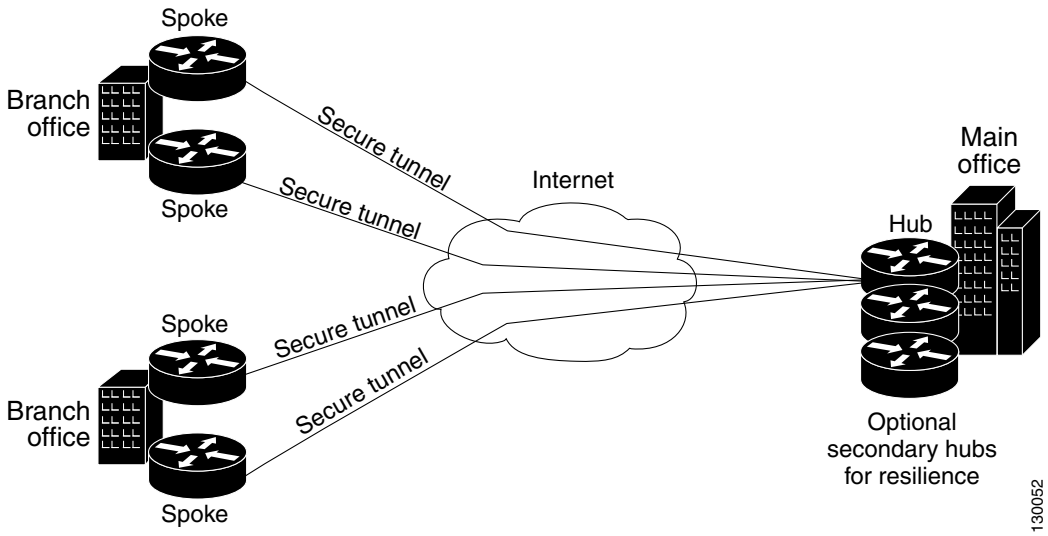

This topology usually represents an intranet VPN that connects an enterprise's main office with branch offices using persistent connections to a third-party network or the Internet. VPNs in a hub-and-spoke topology provide all employees with full access to the enterprise network, regardless of the size, number, or location of its remote operations.

A hub is generally located at an enterprise's main office. Spoke devices are generally located at an enterprise's branch offices. In a hub-and-spoke topology, most traffic is initiated by hosts at the spoke site, but some traffic might be initiated from the central site to the spokes.

If the hub in a hub-and-spoke configuration becomes unavailable for any reason, IPsec failover transfers tunnel connections seamlessly to a failover (backup) hub, which is used by all spokes. You can configure multiple failover hubs for a single primary hub.

In a hub-and-spoke VPN topology, all IPsec technology types can be assigned except GET VPN.

Related Topics

- Understanding IPsec Technologies and Policies, page 25-5
- Implicitly Supported Topologies, page 25-5
- Creating or Editing VPN Topologies, page 25-28
- Chapter 26, "Configuring IKE and IPsec Policies"

Point-to-Point VPN Topologies

In a point-to-point VPN topology, two devices communicate directly with each other, without the option of IPsec failover as in a hub-and-spoke configuration. To establish a point-to-point VPN topology, you specify two endpoints as peer devices. Because either of the two devices can initiate the connection, the assigned IPsec technology type can be only regular IPsec or IPsec/GRE.

In Security Manager, you can configure a special type of regular IPsec point-to-point VPN called an Extranet. An Extranet VPN is a connection between a device in your managed network and an unmanaged device, such as a router in your service provider's network, a non-Cisco device, or simply a device in your network that is being managed by a different group (that is, one that does not appear in the Security Manager inventory).

The following illustration shows a typical point-to-point VPN topology.

Figure 25-2 Point-to-Point VPN Topology

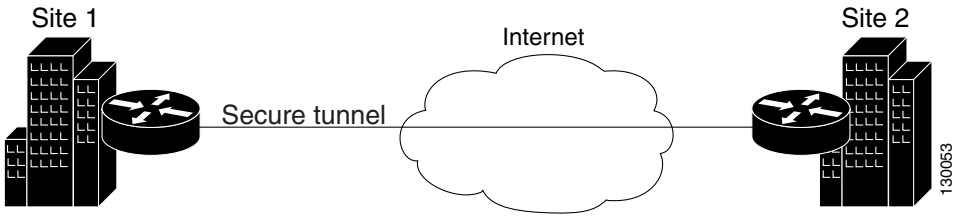

Related Topics

- Understanding IPsec Technologies and Policies, page 25-5
- Implicitly Supported Topologies, page 25-5
- Creating or Editing VPN Topologies, page 25-28
- Creating or Editing Extranet VPNs, page 25-66
- Chapter 26, "Configuring IKE and IPsec Policies"

Full Mesh VPN Topologies

A full mesh topology works well in a complicated network where all peers need to communicate with each other. In this topology type, every device in the network communicates with every other device through a unique IPsec tunnel. All devices have direct peer relationships with one another, preventing a bottleneck at the VPN gateway device, and saving the overhead of encryption and decryption on the device.

You can assign only Regular IPsec, IPsec/GRE, and GET VPN technologies to a full mesh VPN topology.

The following illustration shows a typical full mesh VPN topology.

Figure 25-3 Full Mesh VPN Topology

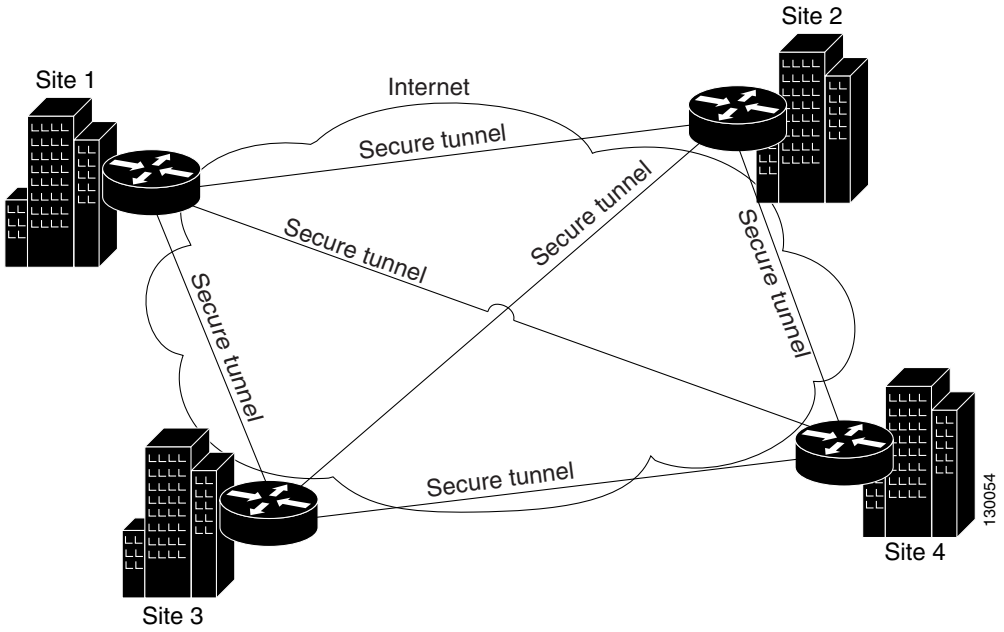

A full mesh network is reliable and offers redundancy. When the assigned technology is GRE and one device (or node) can no longer operate, all the rest can still communicate with one another, directly or through one or more intermediate nodes. With regular IPsec, if one device can no longer operate, a crypto access control list (ACL) that specifies the protected networks, is created per two peers.

GET VPN is based on the group trust model. In this model, group members register with a key server. The key server uses the Group Domain of Interpretation (GDOI) protocol for distributing the security policy and keys for encrypting traffic between the group members. Because you can configure a primary key server and secondary key servers that synchronize their policies with the primary one, if the primary key server becomes unavailable, a secondary key server can take over.

Note When the number of nodes in a full mesh topology increases, scalability may become an issue—the limiting factor being the number of tunnels that the devices can support at a reasonable CPU utilization.

Related Topics

- Understanding IPsec Technologies and Policies, page 25-5
- Implicitly Supported Topologies, page 25-5
- Creating or Editing VPN Topologies, page 25-28
- Chapter 26, "Configuring IKE and IPsec Policies"

Implicitly Supported Topologies

In addition to the three main VPN topologies, other more complex topologies can be created as combinations of these topologies. They include:

- Partial mesh—A network in which some devices are organized in a full mesh topology, and other devices form either a hub-and-spoke or a point-to-point connection to some of the fully meshed devices. A partial mesh does not provide the level of redundancy of a full mesh topology, but it is less expensive to implement. Partial mesh topologies are generally used in peripheral networks that connect to a fully meshed backbone.
- Tiered hub-and-spoke—A network of hub-and-spoke topologies in which a device can behave as a hub in one or more topologies and a spoke in other topologies. Traffic is permitted from spoke groups to their most immediate hub.
- Joined hub-and-spoke—A combination of two topologies (hub-and-spoke, point-to-point, or full mesh) that connect to form a point-to-point tunnel. For example, a joined hub-and-spoke topology could comprise two hub-and-spoke topologies, with the hubs acting as peer devices in a point-to-point topology.

Related Topics

- Creating or Editing VPN Topologies, page 25-28
- Hub-and-Spoke VPN Topologies, page 25-2
- Point-to-Point VPN Topologies, page 25-3
- Full Mesh VPN Topologies, page 25-4

Understanding IPsec Technologies and Policies

Security Manager provides seven types of IPsec technologies that you can configure on the devices in your site-to-site VPN topology—Regular IPsec, IPsec/GRE, GRE Dynamic IP, standard and large scale DMVPN, Easy VPN, and GET VPN. The assigned technology determines which policies you can configure for the VPN.

You assign an IPsec technology to a VPN topology during its creation. After an IPsec technology is assigned to a VPN topology, you cannot change the technology, other than by deleting the VPN topology and creating a new one. See Defining the Name and IPsec Technology of a VPN Topology, page 25-30.

The following topics explain some basic concepts about IPsec technologies and site-to-site VPN policies:

- Understanding Mandatory and Optional Policies for Site-to-Site VPNs, page 25-6
- Overview of Site-to-Site VPN Policies, page 25-8
- Understanding Devices Supported by Each IPsec Technology, page 25-9

- Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11
- Understanding and Configuring VPN Default Policies, page 25-12
- Using Device Overrides to Customize VPN Policies, page 25-13
- Understanding VRF-Aware IPsec, page 25-14

Understanding Mandatory and Optional Policies for Site-to-Site VPNs

Some site-to-site VPN policies are mandatory, which means that you must configure them to create a VPN topology or to save your changes when editing them. Most mandatory policies have predefined defaults, which you can use to complete the definition of a VPN topology, but you typically must edit the policies to ensure their settings work for your network.

Optional policies, which you need to configure only if you desire the services defined by the policy, do not come with predefined defaults.

Tip You can configure your own mandatory policy defaults by creating shared policies that specify the desired settings, and then by selecting these shared policies when creating a VPN. You can even make the shared policies the defaults for the Create VPN wizard. However, these default policies do not apply when you create Extranet VPNs; with Extranet VPNs, you must always configure the settings for mandatory policies as part of the normal wizard flow. In addition, you cannot create a default policy for IKEv2 Authentication. For more information, see Understanding and Configuring VPN Default Policies, page 25-12.

Some mandatory policies are mandatory only under certain conditions. For example, an IKEv1 preshared key policy is mandatory only if the default (mandatory) IKEv1 proposal uses preshared key authentication. If the selected IKE authentication method is Certificate (RSA Signature), an IKEv1 Public Key Infrastructure policy is mandatory (see Deciding Which Authentication Method to Use, page 26-8). If you allow IKEv2 negotiations in the topology, an IKEv2 Authentication policy is mandatory.

The following table lists the mandatory and optional policies for each predefined technology that you can assign to the devices in your site-to-site VPN topology.

Table 25-1 Site-to-Site VPN IPsec Technologies and Policies

Technology	Mandatory Policies	Optional Policies
Regular IPsec See Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19.	• IKE Proposal • IPsec Proposal • When allowing IKEv1, one of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • When allowing IKEv2, IKEv2 Authentication	• VPN Global Settings

Table 25-1 Site-to-Site VPN IPsec Technologies and Policies (continued)

Technology	Mandatory Policies	Optional Policies
IPsec/GRE (Generic Routing Encapsulation) See Understanding GRE, page 27-2.	• IKE Proposal • IPsec Proposal • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • GRE Modes	• VPN Global Settings
GRE Dynamic IP See Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5.	• IKE Proposal • IPsec Proposal • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • GRE Modes	• VPN Global Settings
Dynamic Multipoint VPN (DMVPN). See Understanding DMVPN, page 27-10.	• IKE Proposal • IPsec Proposal • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • GRE Modes	• VPN Global Settings
Large Scale DMVPN See Configuring Large Scale DMVPNs, page 27-16.	• IKE Proposal • IPsec Proposal • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • GRE Modes • Server Load Balance	• VPN Global Settings
Easy VPN See Understanding Easy VPN, page 28-1.	• IKE Proposal • Easy VPN IPsec Proposal • Client Connection Characteristics • If any servers are IOS or PIX 6.3 devices: User Group • If any servers are ASA or PIX 7.0+ devices: Connection Profiles	• IKEv1 Public Key Infrastructure (mandatory if using certificates) • VPN Global Settings

Table 25-1 Site-to-Site VPN IPsec Technologies and Policies (continued)

Technology	Mandatory Policies	Optional Policies
GET VPN See Understanding Group Encrypted Transport (GET) VPNs, page 29-2.	• Group Encryption • IKE Proposal for GET VPN • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure	• Global Settings for GET VPN
Regular IPSec VTI See Configuring Tunnel Interface, page 46-22.	• IKE Proposal • Peers • One of: IKEv1 Preshared Key or IKEv1 Public Key Infrastructure • IKEv2 Authentication • Tunnel Interface with IPSec Profile	

Related Topics

- Creating or Editing VPN Topologies, page 25-28
- Understanding Devices Supported by Each IPsec Technology, page 25-9
- Understanding and Configuring VPN Default Policies, page 25-12
- Understanding VPN Topologies, page 25-2
- Chapter 26, "Configuring IKE and IPsec Policies"
- Understanding Policies, page 5-1

Overview of Site-to-Site VPN Policies

You can access site-to-site VPN policies by selecting **Manage > Site-To-Site VPNs**, or by clicking the **Site-To-Site VPN Manager** button on the toolbar, and then selecting the required policy in the Policies selector of the Site-to-Site VPN window. You can also access site-to-site VPN policies from Device view or Policy view. For more information, see Accessing Site-to-Site VPN Topologies and Policies, page 25-17.

The following is a summary of all of the site-to-site VPN policies, some of which you cannot create as shared policies. Note that some of these policies are documented in the sections that explain remote access VPNs, because the policies are used for both remote access and site-to-site VPNs. However, you must configure these policies separately for each type of VPN.

- Client Connection Characteristics. See Configuring Client Connection Characteristics for Easy VPN, page 28-7.
- Connection Profiles. See Connection Profiles Page, page 31-8.
- Easy VPN IPsec Proposal. See Configuring an IPsec Proposal for Easy VPN, page 28-10.
- GRE Modes. See Understanding the GRE Modes Page, page 27-1.
- Group Encryption Policy. See Defining GET VPN Group Encryption, page 25-54.

- Group Members. See Configuring GET VPN Group Members, page 29-20.
- IKE Proposal. See Configuring an IKE Proposal, page 26-9.
- IKE Proposal for GET VPN. See Configuring the IKE Proposal for GET VPN, page 29-15.
- IKEv2 Authentication. See Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.
- IPsec Proposal. See Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22
- Key Servers. See Configuring GET VPN Key Servers, page 29-18.
- Peers. See Defining the Endpoints and Protected Networks, page 25-34.
- IKEv1 Preshared Key. See Configuring IKEv1 Preshared Key Policies, page 26-48.
- IKEv1 Public Key Infrastructure. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.
- Server Load Balance. See Configuring Server Load Balancing in Large Scale DMVPN, page 27-17.
- User Group Policy. See Configuring a User Group Policy for Easy VPN, page 28-14.
- VPN Global Settings. See Configuring VPN Global Settings, page 26-30.
- Global Settings for GET VPN. See Configuring Global Settings for GET VPN, page 29-16.

Understanding Devices Supported by Each IPsec Technology

Each IPsec technology supports different devices as members of their topology. The following table describes the basic device support. These requirements are enforced when you select devices for a VPN; in some cases, the device lists are filtered to show only supported devices. In other cases, a device might be supported in one role (for example, as a spoke), but not supported in another role; in these cases, you can select the wrong device type, but you are prevented from saving the change (a message will explain the specific problem).

Tip Some device models have NO-VPN versions, which do not support VPN configuration. Thus, although the 3845 model might be supported for a type of VPN, the 3845 NOVPN model is not supported. In addition, the Cisco Catalyst 6500 series ASA Services Module (running software release 8.5(x)) does not support any type of VPN.

Table 25-2 Devices Supported by Each IPsec Technology

Technology	Supported Platforms
Regular IPsec See Chapter 26, "Configuring IKE and IPsec Policies".	Regular IPsec policies can be configured on Cisco IOS security routers (including Aggregation Service Routers, or ASRs), PIX Firewalls, and ASA 5500 series devices. Except for Extranet VPNs, Catalyst VPN service modules are also supported. IKEv2 is supported on ASA release 8.4(x) only. If you limit the topology to IKEv2 only, all devices must support IKEv2. If you allow both IKEv1 and IKEv2, devices that do not support IKEv2 automatically use IKEv1.

Table 25-2 Devices Supported by Each IPsec Technology (continued)

Technology	Supported Platforms
IPsec/GRE (Generic Routing Encapsulation). See Understanding GRE, page 27-2.	GRE policies can be configured on Cisco IOS security routers (including ASRs) and Catalyst 6500/7600 devices.
GRE Dynamic IP. See Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5.	GRE Dynamic IP can be configured on Cisco IOS security routers (including ASRs) and Catalyst 6500/7600 devices.
Dynamic Multipoint VPN (DMVPN), Large Scale DMVPN. See Dynamic Multipoint VPNs (DMVPN), page 27-9 and Configuring Large Scale DMVPNs, page 27-16.	DMVPN configuration is supported on Cisco IOS 12.3T devices and later, and on ASRs running Cisco IOS XE Software 2.x or later (known as 12.2(33)XNA+ in Security Manager). Large Scale DMVPN configuration also supports Catalyst 6500/7600 devices as IPsec Terminators. To use DMVPN phase 3 connections between spokes, devices must run IOS Software release 12.4(6)T or higher; ASRs must run IOS XE Software release 2.4 (called 12.2(33)XND) or higher.
Easy VPN. See Chapter 28, "Easy VPN".	The Easy VPN Server can be a Cisco IOS security router (including ASRs), a Catalyst 6500/7600 (with supported VPN service modules or port adapters), a PIX Firewall, or an ASA 5500 series device. The Easy VPN client is supported on PIX 501, 506, 506E Firewalls running PIX 6.3, Cisco 800-3900 Series routers, and ASA 5505 devices running OS version 7.2 or later.
GET VPN. See Chapter 29, "Group Encrypted Transport (GET) VPNs".	Key servers can be configured on: • Cisco 1800, 2800, 3800 Series ISR, Cisco 7200 Series Routers, and Cisco 7301 Routers running Cisco IOS Software release 12.4(15)T or later. • Cisco 1900, 2900, 3900 Series ISR running release 15.0 or later. Group members can be configured on Cisco 1800, 1900, 2800, 2900, 3800, 3900 Series ISR, Cisco 7200 Series Routers, and Cisco 7301 Routers with the same minimum software releases. The Cisco 871 ISR can also be used as a group member if GET VPN is deployed with very few (1-3) IPSec SAs. In addition, you can configure Cisco ASR Routers using Cisco IOS XE Software Release 2.3 (12.2(33)XNC) and above as group members.

Related Topics

- Creating or Editing VPN Topologies, page 25-28
- Understanding Mandatory and Optional Policies for Site-to-Site VPNs, page 25-6
- Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11
- Understanding and Configuring VPN Default Policies, page 25-12

- Understanding VPN Topologies, page 25-2
- Chapter 26, "Configuring IKE and IPsec Policies"
- Understanding Policies, page 5-1

Including Unmanaged or Non-Cisco Devices in a VPN

Your VPN might include devices that you cannot, or should not, manage in Security Manager. These include:

- Cisco devices that Security Manager supports, but for which your organization is not responsible. For example, you might have a VPN that includes spokes in networks managed by other organizations within your company, or a connection to a service provider or partner network.
- Non-Cisco devices. You cannot use Security Manager to create and deploy configurations to non-Cisco devices.

You have two options for handling these kinds of devices:

- If the connection is a regular IPsec point-to-point connection, you can configure the connection as an Extranet VPN as described in Creating or Editing Extranet VPNs, page 25-66.
- For other types of connections, you can include these devices in the Security Manager inventory as "unmanaged" devices. These devices can serve as endpoints in a VPN topology, but Security Manager does not discover any configurations from the device, nor does it deploy configurations to them.

When the Extranet VPN option will not work, you must do the following before you can add an unmanaged device to a VPN topology:

- Manually add the device as an unmanaged device to the device inventory using the procedure described in Adding Devices by Manual Definition, page 3-26. Ensure that you make the following selections:
  - Select a Cisco device type that is comparable to the device you are adding in terms of VPN-supported technologies. The device type controls the types of VPN topologies to which you can add the device. For example, for GRE/DMVPN, you might select an integrated services router such as an 1800 or 2800 series, whereas in Easy VPN you could also select an ASA or PIX device if appropriate.
  - Deselect the **Manage in Cisco Security Manager** option. This is very important, because the default is to make all new devices managed devices. If you forget to do this while adding the device, you can deselect the option later on the General tab in the device properties (right-click the device and select **Device Properties**).
- Using the interface policy for the device, define the external VPN interface to which managed devices will point. Because the device is unmanaged, your definitions in this policy are never configured on the device; the policy simply represents what you have configured on the device outside of Security Manager.

Related Topics

- Understanding Devices Supported by Each IPsec Technology, page 25-9
- Selecting Devices for Your VPN Topology, page 25-32
- Creating or Editing VPN Topologies, page 25-28

Understanding and Configuring VPN Default Policies

For most VPN policies that are mandatory, Security Manager includes "factory default" settings for the policies. These defaults are generic, and might not be appropriate for your network, but they do allow you to complete the creation of a VPN without having to stop and start over when you do not have the needed shared policy configured. Therefore, you can, and should, create your own default VPN policies for mandatory policies. You can also create defaults for certain optional policies.

Before configuring new defaults, consider the types of VPNs you are likely to configure, then review the types of policies for which you can create defaults. Select **Tools > Security Manager Administration**, then select **VPN Policy Defaults** from the table of contents. Select the tabs for the desired IPsec technologies to see which policies are available. If a policy is assigned Factory Default, or if this option is available from the drop-down list, the policy is mandatory; other policies are optional. You can also create default policies for remote access VPNs, and for site-to-site endpoint configurations. Click the **View Content** button next to a selected policy to see the policy definition.

The following procedure explains how to create and use VPN policy defaults.

Tips

- When you configure VPN default policies, you are selecting shared policies. Although you can configure only one default per policy per IPsec technology, users can select different shared policies when configuring VPNs. Thus, you might want to configure more than one shared policy that users can select, and configure the most commonly-used policy as the default policy. For more information about how users can select different policies when configuring a VPN, see Assigning Initial Policies (Defaults) to a New VPN Topology, page 25-62.
- Although the IKEv2 Authentication policy is a mandatory policy for topologies that allow IKEv2 negotiations, there are no IKEv2 Authentication factory default settings, and you cannot create IKEv2 Authentication shared policies. Therefore, whenever you allow IKEv2 in a topology, you must manually configure the IKEv2 Authentication policy before the topology is valid.
- The Public Key Infrastructure policy is required for IKEv1 if you configure the IKE Proposal policy to use certificate authentication. However, there are no factory default settings for this policy, so if you intend to use certificate authentication for IKEv1, consider creating default Public Key Infrastructure policies.
- Keep in mind that any change to a shared policy affects all VPNs that are using the policy. This can make it easy to implement across-the-board changes that are required for every VPN. However, after creating the VPN, the user can switch from a shared policy to a local policy, so that any changes to the configuration must be done specifically for the VPN topology. For more information about shared policies, see Managing Shared Policies in Policy View, page 5-50.
- These default policies do not apply when you create Extranet VPNs. With Extranet VPNs, you must always configure the settings for mandatory policies as part of the normal wizard flow.

Step 1 Create the default policies. All default policies are shared policies.

a. In Policy view (select **View > Policy View**), select the policy for which you want to configure defaults. The policies are in the **Site-to-Site VPN** or **Remote Access VPN** folders.

b. Click the **Create a Policy (+)** button at the bottom of the shared policy selector, enter a name for the policy, and click **OK**.

c. Configure the desired settings. Click the **Help (?)** button in the toolbar to get reference information about the settings available in the selected policy.

d. Repeat the process until you have created at least one shared policy for each policy for which you want to define a default policy.

Step 2 If desired, create defaults for the VPN endpoints. These defaults are interface role objects, which identify the interface names used for VPN connections (for example, GigabitEthernet0/1). Create separate roles for internal and external VPN interfaces.

a. Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

b. Select **Interface Roles** from the table of contents.

c. Click the **New Object (+)** button, enter the interface name patterns that identify the most commonly used interfaces for VPN internal or external interfaces in your network, and click **OK**.

For more information about interface roles and the wildcards you can use to configure them, see Understanding Interface Role Objects, page 6-73 and Creating Interface Role Objects, page 6-74.

Step 3 Submit the policies and policy objects to the database. You will have to resolve any validation errors.

- In non-Workflow mode, select **File > Submit**.
- In Workflow mode without an activity approver, select **Activities > Approve Activity**.
- In Workflow mode with an activity approver, select **Activities > Submit Activity**. You will have to wait for the activity to be approved before you can select the policies and objects as defaults.

Step 4 Select your newly-configured policies and policy objects as VPN policy defaults.

a. Select **Tools > Security Manager Administration**, and then select **VPN Policy Defaults** from the table of contents (see VPN Policy Defaults Page, page 11-74).

b. Select the desired tabs, then select the policies you configured from the drop-down lists for each of the mandatory or optional policies for which you configured defaults.

On the S2S Endpoints tab, select the appropriate interface role objects.

c. Click **Save** to save your defaults.

The next time a user runs the Create VPN wizard, the defaults you selected will be used as the wizard's defaults. Users can select any other shared policy or interface role to override the default.

Using Device Overrides to Customize VPN Policies

Many VPN policies use Security Manager policy objects in their configuration. Policy objects are containers that allow you to create reusable configurations.

Because a VPN policy applies to every device in a VPN topology, you might need to make modifications to a policy object used in a policy for certain devices within the VPN topology. There might even be situations where you need to make modifications for all devices within a topology. You accomplish these modifications with device-level overrides on the policy objects.

For example, when defining a PKI policy, you need to select a PKI enrollment object. If the hub of your VPN uses a different CA server than the spokes, you must use device-level overrides to specify the CA server used by the hub. Although the PKI policy references a single PKI enrollment object, the actual CA server represented by this object differs for the hub, based on the device-level override you define.

To enable a policy object to be overridden, you must select the **Allow Override per Device** option in the policy object definition. You can then create device-level overrides. For more information about overriding a VPN policy object at the device level, see the following topics:

- Understanding Policy Object Overrides for Individual Devices, page 6-18
- Allowing a Policy Object to Be Overridden, page 6-18
- Creating or Editing Object Overrides for a Single Device, page 6-19

- Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19

Understanding VRF-Aware IPsec

One obstacle to successfully deploying peer-to-peer VPNs is the separation of routing tables, and the use of overlapping addresses, which usually results from using private IP addresses in customer networks. The VRF-Aware IPsec feature, which introduces IPsec tunnel mapping to Multiprotocol Label Switching (MPLS) VPNs, solves this problem.

The VRF-Aware IPsec feature enables you to map IPsec tunnels to Virtual Routing Forwarding (VRF) instances, using a single public-facing address. A VRF instance defines the VPN membership of a customer site attached to the Provider Edge (PE) router. A VRF comprises an IP routing table, a derived Cisco Express Forwarding (CEF) table, a set of interfaces that use the forwarding table, and a set of rules and routing protocol parameters that control the information that is included in the routing table. A set of routing and CEF tables is maintained for each VPN customer across the MPLS/VPN network.

Since each VPN has its own routing and forwarding table in the router, any customer or site that belongs to a VPN is provided access only to the set of routes contained within that table. Any PE router maintains a number of routing tables and a global routing table per VPN, which can be used to reach other routers in the provider network. Effectively, a number of virtual routers are created in a single physical router. Across the MPLS core to the other PE routers, this routing separation is maintained by adding unique VPN identifiers, such as the route distinguisher (RD).

Note VRF-Aware IPsec can also be configured on devices in a remote access VPN. For more information, see Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices), page 33-7.

In Security Manager, you can configure VRF-Aware IPsec in your hub-and spoke VPN topology, with either a single device providing all functionality (“one-box” solution) or with multiple devices, each providing a part of the functionality (“two-box” solution). The solution of one device providing all the functionality can affect performance by overloading the system, whereas separating the functionality in a two-box solution provides better scaling for each function.

The following topics describe:

- VRF-Aware IPsec One-Box Solution, page 25-14
- VRF-Aware IPsec Two-Box Solution, page 25-15
- Enabling and Disabling VRF on Catalyst Switches and 7600 Devices, page 25-17

For information on configuring VRF-aware IPsec, see Configuring VRF Aware IPsec Settings, page 25-48.

VRF-Aware IPsec One-Box Solution

In the one-box solution, IPsec tunnels terminate on a Cisco IOS router, which serves as the Provider Edge (PE) device. The PE device maps these tunnels to the appropriate MPLS/VPN network and serves as the IPsec Aggregator, by performing IPsec encryption and decryption from the Customer Edge (CE) devices.

Note The configuration of routing between the PE device and the MPLS cloud is done by Cisco IP Solution Center. See the *Cisco IP Solution Center MPLS VPN User Guide*.

The following illustration shows the topology of a one-box solution.

Figure 25-4 VRF-Aware IPsec One-Box Solution

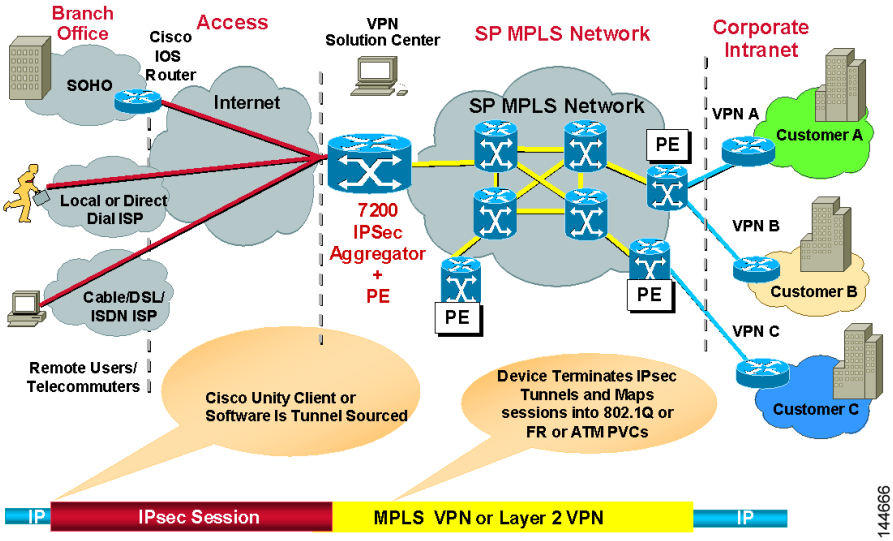

Related Topics

- Understanding VRF-Aware IPsec, page 25-14
- Configuring VRF Aware IPsec Settings, page 25-48
- Defining the Endpoints and Protected Networks, page 25-34

VRF-Aware IPsec Two-Box Solution

In the two-box solution, the PE device does just the MPLS mapping, while a separate IPsec Aggregator device does the IPsec encryption and decryption from the CEs.

Note Security Manager fully manages the IPsec Aggregator, including routing to the PE device. The PE device is fully managed by Cisco IP Solution Center. This includes routing between the PE device and the MPLS cloud, and routing from the PE to the IPsec Aggregator. For more information, see the *Cisco IP Solution Center MPLS VPN User Guide*.

The following illustration shows the topology of a two-box solution.

Figure 25-5 VRF-Aware IPsec Two-Box Solution

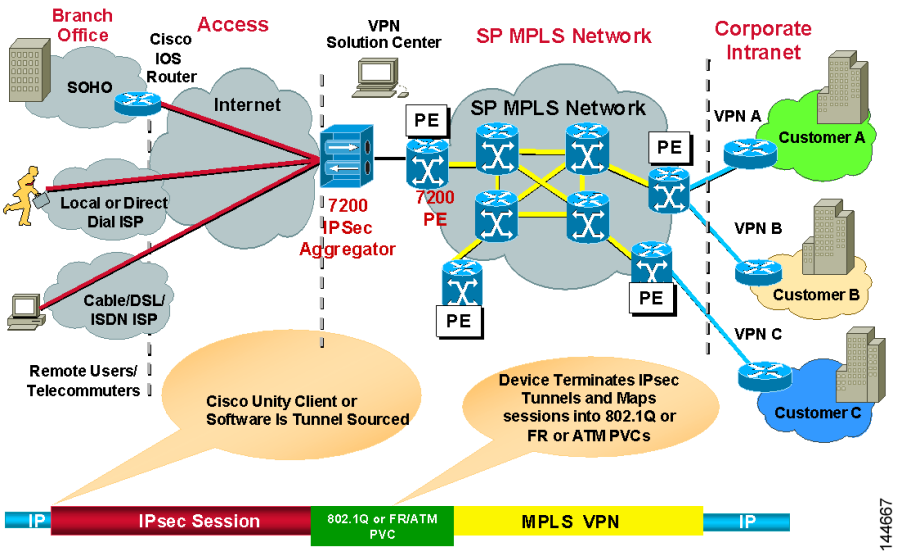

Using the two-box solution, you configure VRF-Aware IPsec on devices in your VPN topology, as follows:

1. Configure the connection between the IPsec Aggregator and the PE device.

   Create a hub-and-spoke VPN topology and assign an IPsec technology to it. In this topology, the hub is the IPsec Aggregator, and the spokes may be Cisco IOS routers, PIX Firewalls, Catalyst VPN service modules, or Adaptive Security Appliance (ASA) devices. The IPsec Aggregator may be a security router or a Catalyst VPN service module. You then define the VRF parameters (VRF name and unique routing identifier) on the hub.

   Note VRF-Aware IPsec supports the configuration of IPsec, GRE, or Easy VPN technologies on Cisco IOS routers and Catalyst VPN service modules. DMVPN is also supported, but only on Cisco IOS routers.

2. Specify the VRF forwarding interface (or VLAN for a Catalyst VPN service module) between the IPsec Aggregator and the PE device.
3. Define the routing protocol and autonomous system (AS) number to be used between the IPsec Aggregator and the PE. Available routing protocols include BGP, EIGRP, OSPF, RIPv2, and Static Route.

   If the routing protocol defined between the IPsec Aggregator and the PE differs from the routing protocol used for the secured IGP, routing is redistributed to the secured IGP, using this routing protocol and AS number. Routing is also redistributed from the secured IGP to the PE.

   Note Redistributing the routing is only relevant when IPsec/GRE or DMVPN is the selected technology.

Related Topics

- Understanding VRF-Aware IPsec, page 25-14
- Configuring VRF Aware IPsec Settings, page 25-48

- Defining the Endpoints and Protected Networks, page 25-34

Enabling and Disabling VRF on Catalyst Switches and 7600 Devices

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, it does not support any enhancements as Cisco Catalyst switches is now End of Life.

Deployment fails when you change the virtual routing and forwarding (VRF) mode on the Catalyst switches and 7600 hub of an existing site-to-site VPN. For example, if you initially configured VRF in the Create VPN wizard and deployed, but later return to the Peers policy and deselect the Enable VRF Settings check box, deployment fails. (This setting is found in the VRF Aware IPSec tab of the Edit Endpoints dialog box; see Configuring VRF Aware IPsec Settings, page 25-48.) Deployment likewise fails if you try to enable VRF on a VPN that was not initially configured with it.

You cannot change the VRF mode on a Catalyst 6500/7600 during VPN operation. This restriction applies only to Catalyst 6500/7600 hubs, not to any other device type.

This restriction does not apply to changes made to the VRF settings themselves. For example, if VRF is configured on the VPN topology, you can return to the Peers policy and change the VRF name or route distinguisher.

If you need to change the VRF mode of a VPN, and you are using Catalyst 6500/7600 devices as hubs, use the following procedure.

Related topics

- Understanding VRF-Aware IPsec, page 25-14
- VRF-Aware IPsec One-Box Solution, page 25-14
- VRF-Aware IPsec Two-Box Solution, page 25-15

Step 1 Delete the VPN topology from Security Manager.

Step 2 Deploy your changes.

Step 3 Reload (restart) the Catalyst 6500/7600 device.

Step 4 Right-click the device in Security Manager and select **Discover Policies on Device**. Perform a complete policy rediscovery.

Step 5 Open the Create VPN wizard and redefine the VPN topology. At this point, you can select a different VRF mode. See Configuring VRF Aware IPsec Settings, page 25-48 and Creating or Editing VPN Topologies, page 25-28.

Accessing Site-to-Site VPN Topologies and Policies

You can use the following methods to access and configure site-to-site VPN topologies and policies:

- **Site-to-Site VPN Manager**—This is the main tool for configuring VPN topologies. You can view a list of all site-to-site VPNs configured in Security Manager and edit their configurations and policies, including device membership. For information on using this tool, see Site-to-Site VPN Manager Window, page 25-18.

- **Site-to-Site VPN policy in Device view**—When you select a device in device view, you can select the Site-to-Site VPN policy in the Policies selector to see a list of all site-to-site VPNs in which the device participates and edit those topologies. You can also create new VPNs, or select a VPN and open the Site-to-Site VPN Manager to edit the policies for the selected VPN. This device view policy is essentially a short-cut into the Site-to-Site VPN Manager. For more information about using this policy, see Configuring VPN Topologies in Device View, page 25-19.
- **Site-to-Site VPN folder in Policy view**—Policy view is used to create shared policies. Many of the site-to-site VPN policies are shareable. Thus, you can configure shared policies that you can assign to more than one VPN topology while configuring the topology in the Site-to-Site VPN Manager. You can configure shared policies as defaults for the Create VPN wizard, as described in Understanding and Configuring VPN Default Policies, page 25-12.

  You can also create shared policies from the Site-to-Site VPN Manager window in much the same way you can create them from local policies in Device view, although all sharing commands in the Site-to-Site VPN Manager window are available only on the right-click context menu (when right-clicking a shareable policy).

  For more information on creating shared policies in Policy view, see Managing Shared Policies in Policy View, page 5-50.

Site-to-Site VPN Manager Window

The Site-to-Site VPN Manager lists all site-to-site VPNs configured in Security Manager. The VPNs selector, in the upper left pane of the window, lists all existing VPN topologies (see Understanding VPN Topologies, page 25-2). An icon indicates the type of VPN (hub and spoke, point to point, or full mesh).To view or edit a topology, select it, and its policies are loaded into the policy selector in the lower left pane. Select a policy to see its definition in the right pane.

To open the Site-to-Site VPN Manager, click the **Site-To-Site VPN Manager** button on the toolbar or select **Manage > Site-To-Site VPNs**.

Use the Site-to-Site VPN Manager window to:

- Create, edit, and delete VPN topologies.
  - To create a VPN topology, click the **Create VPN Topology (+)** button above the VPN selector and select the type of topology you want to create from the options that are displayed. This action opens the Create VPN Wizard or the Create Extranet VPN wizard. For more information, see Creating or Editing VPN Topologies, page 25-28 or Creating or Editing Extranet VPNs, page 25-66.
  - To edit a VPN topology, select it and click the **Edit VPN Topology (pencil)** button, or right-click it and select **Edit**. This opens the Edit VPN or Edit Extranet VPN dialog box, which contains the most of the same pages as the Create VPN wizard in a tabbed layout.
  - To delete a VPN topology, select it and click the **Delete VPN Topology (trash can)** icon, or right-click it and select **Delete**. You are asked to confirm the deletion. See Deleting a VPN Topology, page 25-71.
- View detailed information about each VPN topology; select the topology, then select the VPN Summary policy. See Viewing a Summary of a VPN Topology's Configuration, page 25-63.
- View and configure the endpoints defined for a VPN topology. You can see endpoints on the Endpoints tab or when editing a VPN topology, or by selecting the **Peers** policy. For GET VPN topologies, there is no Peers policy; instead, use the **Key Servers** and **Group Members** policies to view and configure endpoints. For Extranet VPNs, the endpoints are on the Device Selection tab when editing the VPN, or also in the Peers policy.

- View and edit the policies assigned to a VPN topology, assign shared policies, or create shared policies from existing policies. For information on individual policies, see Overview of Site-to-Site VPN Policies, page 25-8.

  The options and methods for configuring shared policies from the Site-to-Site VPN Manager are the same as those from Device view, as explained in the sections under Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37 and Using the Policy Banner, page 5-38. You can share, assign, unassign, edit assignments, and rename policies, but no VPN policies allow inheritance. To perform these tasks, select the VPN topology, then right-click the desired policy and select the desired command.

  You can also use Policy view to configure shared VPN policies.

Configuring VPN Topologies in Device View

Use the Site-to-Site VPN Device view policy to view and edit the site-to-site VPN topologies to which a device belongs, if any. You can edit the VPN policies and change whether the device participates in the topology. You can also create new VPN topologies.

This policy is essentially an access point for the Site-to-Site VPN Manager (see Site-to-Site VPN Manager Window, page 25-18).

To open this policy, in Device view, select the desired device and then select **Site-to-Site VPN** from the Policy selector.

The VPN topologies table lists all of the site-to-site VPNs to which this device belongs. Information includes the type of VPN, its name, IPSec technology, and description. Beginning with version 4.9, Security Manager also displays the Last Modified Ticket information for the VPN topologies. The VPN topologies that have been created or edited using the Ticket Management system will have the last Modified Ticket ID information available on this page. You can also filter the VPN topologies by the Last Modified Ticket ID.

- To add a VPN, click the **Create VPN Topology** button, or right-click in the table and select **Create VPN Topology** and select the type of topology you want to create from the options that are displayed. This action opens the Create VPN Wizard or the Create Extranet VPN wizard. For more information, see Creating or Editing VPN Topologies, page 25-28 or Creating or Editing Extranet VPNs, page 25-66.
- To edit a VPN, select it and click the **Edit VPN Topology** button, right-click the VPN and select **Edit VPN Topology**, or simply double-click the entry. This opens the Edit VPN or Edit Extranet VPN dialog box, which is a tabbed version of the Create VPN wizard (see Creating or Editing VPN Topologies, page 25-28 or Creating or Editing Extranet VPNs, page 25-66).
- To edit the policies for a VPN, select it and click the **Edit VPN Policies** button. The Site-to-Site VPN Window opens displaying information about the VPN topology; select the desired policy from the Policies selector to edit it.
- To delete a VPN, select it and click the **Delete VPN Topology** button, or right-click the VPN and select **Delete VPN Topology**. You are asked to confirm the deletion. For more information, see Deleting a VPN Topology, page 25-71.

Site-To-Site VPN Discovery

You can discover the VPN topologies that are already deployed in your network so that you can use Security Manager to manage them. Your VPN configurations are brought into Security Manager and displayed as site-to-site VPN policies.

Except for Extranet VPNs, you can also rediscover the configurations of existing VPN topologies that are already managed with Security Manager. For information about Site-to-Site VPN rediscovery, see Rediscovering Site-to-Site VPNs, page 25-27.

Note You can also discover configurations on devices in remote access VPNs that are already deployed in your network. See Discovering Remote Access VPN Policies, page 30-12.

These topics provide information about Site-to-Site VPN discovery:

- Supported and Unsupported Technologies and Topologies for VPN Discovery, page 25-20
- Prerequisites for VPN Discovery, page 25-21
- VPN Discovery Rules, page 25-22
- Discovering Site-to-Site VPNs, page 25-24
- Defining or Repairing Discovered VPNs with Multiple Spoke Definitions, page 25-26
- Rediscovering Site-to-Site VPNs, page 25-27

Supported and Unsupported Technologies and Topologies for VPN Discovery

This topic lists the technologies and topologies that Security Manager can discover, as well as the VPN features that are provisioned by Security Manager but cannot be discovered.

Supported Technologies for VPN Discovery

- IPsec, including LAN-to-LAN configurations on ASA devices.
- IPsec + GRE
- IPsec + GRE dynamic IP
- DMVPN
- Easy VPN
- GET VPN

Supported Topologies for VPN Discovery

- Point to point
- Hub and spoke
- Full mesh
- Extranet VPN (point-to-point to an unmanaged device)

VPN Features Provisioned by Security Manager but Unsupported for VPN Discovery

- Large Scale DMVPN with IPsec Terminator (high-concentration hub)
- VRF-Aware IPsec

- Dial backup
- IPsec and ISAKMP profiles for Easy VPN
- Easy VPN with High Availability

If you define and deploy policies of these types using Security Manager, your policies overwrite the device configurations that were not discovered. Therefore, if you want Security Manager to manage existing configurations, you should define policies that match the existing configurations as closely as possible. (Use **Tools > Preview Configuration** to examine the results before deploying.) The VPN provisioning mechanism leverages the content of the existing configuration as much as possible (assuming the content matches the policies configured in Security Manager), but does not retain the naming conventions used in the CLI commands.

Related Topics

- Prerequisites for VPN Discovery, page 25-21
- VPN Discovery Rules, page 25-22
- Discovering Site-to-Site VPNs, page 25-24

Prerequisites for VPN Discovery

For successful VPN discovery, the following prerequisites must be met:

- Except for Extranet VPNs, all devices participating in the VPN must be added to the Security Manager inventory.
- You must provide Security Manager with some basic information about the VPN. The VPN discovery wizard prompts you for the following information:
  - VPN topology (hub and spoke, point to point, full mesh, Extranet).
  - VPN technology (Regular IPsec, IPsec/GRE, GRE dynamic IP, DMVPN, Easy VPN, GET VPN).
  - Devices in the VPN and their roles (hub/spoke). For Extranet VPNs, you specify the managed device only.
  - Source of the VPN configuration. The VPN can be discovered directly from the live network or from Security Manager's Configuration Archive.
- Each device in the VPN must have a crypto map associated with a physical interface. This rule does not apply to the remote (unmanaged) devices in an Extranet VPN.
- If you use OSPF as your routing protocol in a VPN topology, all devices in the VPN must use the same OSPF process number.
- Each PIX 6.3 or ASA 5505 client device in an Easy VPN topology must have a vpnclient configuration.

Related Topics

- Supported and Unsupported Technologies and Topologies for VPN Discovery, page 25-20
- VPN Discovery Rules, page 25-22
- Discovering Site-to-Site VPNs, page 25-24

VPN Discovery Rules

The following table describes the rules by which Security Manager translates and discovers your VPN configurations, and how it handles instances where your configuration on the device does not match what is supported by Security Manager.

Tip Because Extranet VPN discovery involves the analysis of a single device (the managed device), most of these rules do not apply to Extranet VPN discovery. Any rule that involves consistency of values among devices in the VPN is not applicable.

Table 25-3 VPN Discovery Rules

If this condition exists:	The VPN discovery is handled as follows:
Security Manager cannot contact a device in the VPN for live device discovery.	• If the device is the only hub or spoke in the VPN, discovery fails. • If there are other hubs or spokes in the VPN, discovery proceeds but the unavailable device is not discovered. • Except for Extranet VPNs, if the device is a peer in a point-to-point topology, discovery fails. For Extranet VPNs, only the managed device is contacted, and discovery fails if it cannot be contacted. • If the device is a peer in a full mesh topology and there are only two devices, including the unavailable one, in the topology, discovery fails. If there are more than two devices, discovery proceeds but the unavailable device is not discovered.
The VPN is a LAN-to-LAN VPN on an ASA.	The ASA documentation uses “LAN-to-LAN” as a synonym for “site-to-site.” In a LAN-to-LAN VPN configuration, the ASA uses tunnel groups, which when used in a remote access VPN configuration, Security Manager discovers as connection profiles. When discovering site-to-site VPNs on an ASA that uses LAN-to-LAN (L2L) tunnel groups, Security Manager creates a site-to-site VPN topology, and the L2L tunnel groups are not presented to you as connection profiles. Instead, you edit the properties of the VPN topology, and during deployment, Security Manager will translate the configuration into the appropriate L2L tunnel group commands.

Table 25-3 VPN Discovery Rules (continued)

If this condition exists:	The VPN discovery is handled as follows:
There are inconsistencies in the policies or values in the VPN configurations across the devices in the VPN.	• If the values on the hub and the spokes differ, preference is given to the values on the hub. • If a simple selection of one policy or value from several eligible policies or values is required and does not put functionality at risk, Security Manager selects a single policy/value that is common to all devices. For example, a VPN can have a single IKE policy only, whereas there can be more than one IKE policy on the devices. • If selecting one value puts the functionality at risk, no value is discovered for the policy and a validation message is received upon deployment. • If numeric values differ, a message is generated during discovery, and the lower value is discovered. For example, the lowest SA lifetime value in an IPsec policy. • If none of the above options are possible, VPN discovery fails.
Preshared key configuration—there is a different key per set of peers.	The preshared key policy is not discovered; you will have to configure it after discovery is completed. Security Manager discovers preshared key policies only when the preshared key has the same value on all devices.
There is more than one eligible crypto map on the device.	The crypto map that is associated with all or the majority of the devices selected for VPN discovery is used.
A spoke does not have a crypto map associated with the hub.	VPN discovery proceeds but the spoke is not discovered and an error message is generated.
A device does not have the selected transform set value.	VPN discovery proceeds but the device might removed from the VPN topology.
A device does not have the selected IKE proposal.	VPN discovery proceeds but the device might removed from the VPN topology.
A device supports DVTI, but does not have DVTI or a crypto map configured.	VPN discovery fails.
A server supports DVTI, but does not have an IP address configured in the DVTI configuration.	VPN discovery proceeds but with a warning.
A client does not support DVTI.	If the hub is configured with DVTI, discovery proceeds without any warning or Error.
A Hub and Spoke topology where the spokes are not using the same VPNSPA/VSPA slot on the hub (Catalyst 6500/7600).	VPN discovery fails.
The same set of key servers and group members are participating in more than one GET VPN.	Security Manager discovers only one of the topologies.

Table 25-3 VPN Discovery Rules (continued)

If this condition exists:	The VPN discovery is handled as follows:
A User Group policy is configured with backup servers using hostnames instead of an IP addresses.	VPN policy discovery fails with the following error: Policy Discovery Failed: com.cisco.nm.vms.discovery.DiscoveryException: Internal Error In order for discovery to be successful, you need to reconfigure the user group policy on the device with backup servers using IP address, not hostnames.

Related Topics

- Supported and Unsupported Technologies and Topologies for VPN Discovery, page 25-20
- Prerequisites for VPN Discovery, page 25-21
- Discovering Site-to-Site VPNs, page 25-24
- Rediscovering Site-to-Site VPNs, page 25-27

Discovering Site-to-Site VPNs

This procedure describes how to discover a Site-to-Site VPN that is already working in your network but that has not yet been defined in Security Manager.

Related Topics

- Discovering Site-to-Site VPNs, page 25-24
- Discovering Policies, page 5-12
- Supported and Unsupported Technologies and Topologies for VPN Discovery, page 25-20
- Prerequisites for VPN Discovery, page 25-21
- VPN Discovery Rules, page 25-22
- Understanding Devices Supported by Each IPsec Technology, page 25-9
- Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11

Step 1 In Device view, select **Policy > Discover VPN Polices** to open the Discover VPN Policies Wizard—Name and Technology page.

Step 2 Specify the following information:

- **VPN Name**—The name of the VPN being discovered.

  You cannot specify the name when discovering Extranet VPNs. Instead, Security Manager discovers all Extranets defined on the device, and for each Extranet, the VPN name is a hyphenation of the local and remote IP addresses. For example, if the local address is 10.100.10.1 and the remote address is 10.100.11.1, the Extranet VPN is named **10.100.10.1-10.100.11.1**.

- **Description**—An optional description of the VPN. You cannot add a description to Extranet VPN discovery.

- **Topology**—The type of VPN that you are discovering—Hub and Spoke, Point to Point, Full Mesh, or Extranet.

- **IPsec Technology**—The IPsec technology assigned to the VPN—Regular IPsec, IPsec/GRE, GRE Dynamic IP (sub-technology), DMVPN, Easy VPN, GET VPN, or Regular IPSEC VTI. The topology you select controls what is available in this list.

  If you selected IPsec/GRE, you must also specify the type which may be **Standard** (for IPsec/GRE) or **Spokes with Dynamic IP** (to configure GRE Dynamic IP).

Note You can select Regular IPSEC VTI for tunnel based routing as applicable for Hub and Spoke, and Point to Point topologies.

- **Discover From**—You can either discover the VPN directly from the network or from Configuration Archive.
  - Network—Security Manager connects to all live devices to obtain the device configuration. For Extranet VPN discovery, Security Manager connects to the single managed device that you specify.
  - Config Archive—Discovery from Configuration Archive is recommended if you deploy to configuration files instead of live devices. The most recent version of the device configuration in Configuration Archive is used for all devices.

Step 3 Click **Next** to open the Discover VPN Policies Wizard—Device Selection Page.

Step 4 Select the devices participating in the VPN and their role in the VPN (hub, spoke, peer one, peer two, local device, key server, group member, or simply selected devices for full-mesh VPNs) depending on the topology type. For Easy VPN topologies, servers are hubs and clients are spokes.

If there are two or more IPsec terminators in a hub-and-spoke VPN, use the Up and Down arrow buttons to ensure the primary hub is listed first. When there is only one IPsec terminator, regardless of how many hubs are connected to the same IPsec terminator, it is not possible to designate one hub as the primary hub.

For more detailed information on selecting devices for a VPN, see Selecting Devices for Your VPN Topology, page 25-32.

Step 5 Click **Finish** to close the wizard and start the discovery process. The Discovery Status window opens and displays the status of the discovery and indicates whether the discovery of each device has been successful or has failed (see Viewing Policy Discovery Task Status, page 5-22). Error or warning messages are provided to indicate the source of any problems, which may be VPN specific or device specific.

Except for Extranet discovery, when the discovery process completes successfully, and you close the Discovery Status dialog box, the Site-to-Site VPN Manager window opens, displaying summary information for the VPN that was discovered. For Extranet discovery, you must either manually open the Site-to-Site VPN Manager, or select the Site-to-Site VPN policy in Device view, to see the list of discovered Extranet VPNs.

Step 6 Verify that the VPN polices are as required. Edit the policies as necessary.

Tip When discovering Extranet VPNs, all Extranet VPNs defined on the selected device are discovered. Delete the ones that you do not want to manage in Security Manager.

Defining or Repairing Discovered VPNs with Multiple Spoke Definitions

If you discover a VPN whose spokes contain different definitions (for example, different client modes for Easy VPN spokes), Security Manager changes the definitions during discovery to create a uniform definition for all spokes. This behavior occurs because VPN topologies in Security Manager can contain only one set of spoke definitions.

If you want to maintain your original definitions, or create a new VPN that has spokes with different definitions, you can choose one of two approaches:

- Define multiple VPN topologies in Security Manager, where each topology includes spokes containing matching spoke definitions.
- Define a FlexConfig policy that contains the specialized definition, then assign the policy to the spokes that require this definition, as described in the following procedure.

Related Topics

- Creating a New Shared Policy, page 5-54
- Creating FlexConfig Policy Objects, page 7-28
- Modifying Policy Assignments in Policy View, page 5-54
- Site-To-Site VPN Discovery, page 25-20
- Discovering Site-to-Site VPNs, page 25-24
- VPN Discovery Rules, page 25-22

Step 1 Create a shared FlexConfig policy in Policy view:

a. Select **View > Policy View**.

b. Right-click **FlexConfigs** in the Policy Type selector, then select **New FlexConfigs Policy**.

c. Enter a name for the policy and click **OK**.

Step 2 Define the FlexConfig policy by creating and selecting a FlexConfig object:

a. In the work area of Policy view, click the **Add** button on the Details tab.

b. In the FlexConfigs Selector, click the **Create** button in the lower-left corner of the window to open Add or Edit FlexConfig Dialog Box, page 7-30.

c. Define an appended FlexConfig object that contains the required client definition. For example, to define the client mode on an Easy VPN spoke, enter the following commands:

crypto ipsec client ezvpn CSM_EASY_VPN_CLIENT_1

mode client

exit

d. After you create the FlexConfig object, add it to the FlexConfig policy using the selector.

Step 3 In the work area of Policy view, use the Assignments tab to select the spokes to which this policy should be assigned, then click **Save**.

Step 4 Deploy the policy.

Rediscovering Site-to-Site VPNs

You can rediscover the configurations of existing VPN topologies that are already managed with Security Manager so that you do not have to recreate policies changes in the application.

The same rules by which Security Manager translates and discovers VPN configurations apply also to rediscovery. However, you can perform rediscovery only on devices that participate in a VPN topology, and you cannot make any changes to the IPsec technology or topology type. Only the configurations of device specific policies, such as VPN interfaces and protected networks, and any High Availability (HA) policies that are configured on hubs, can be rediscovered. VPN global policies, such as IKE proposals or PKI enrollments, cannot be rediscovered. In addition, you cannot rediscover the following topologies:

- Easy VPN topologies with Dynamic VTI
- Extranet VPNs

This procedure describes how to rediscover the configurations of a Site-to-Site VPN topology that already exists in Security Manager.

Related Topics

- Discovering Site-to-Site VPNs, page 25-24
- Discovering Policies, page 5-12
- Prerequisites for VPN Discovery, page 25-21
- VPN Discovery Rules, page 25-22
- Understanding Devices Supported by Each IPsec Technology, page 25-9
- Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11

Step 1 In the Site-to-Site VPN Manager window, right-click the VPN topology whose configurations you want to rediscover and select **Rediscover Peers**. This opens the Rediscover VPN Policies Wizard—Name and Technology page.

This page displays the type of topology and IPsec technology used in the VPN, which you cannot change.

Step 2 Specify the following information:

- **VPN Discovery Name**—The name of the rediscover VPN job.
- **Description**—An optional description of the VPN.
- **Discover From**—You can either rediscover the VPN directly from the network or from Configuration Archive.
  - Network—Security Manager connects to all live devices to obtain the device configuration.
  - Config Archive—Rediscovery from Configuration Archive is recommended if you deploy to configuration files instead of live devices. The most recent version of the device configuration in Configuration Archive is used for all devices.

Step 3 Click **Next** to open the Rediscover VPN Policies Wizard—Device Selection page.

Step 4 Select the devices whose peer level policies need to be rediscovered and their role in the VPN (hub, spoke, peer one, peer two, key server, group member, or simply selected devices for full-mesh VPNs) depending on the topology type. For Easy VPN topologies, servers are hubs and clients are spokes.

If there are two or more IPsec terminators in a hub-and-spoke VPN, use the Up and Down arrow buttons to ensure the primary hub is listed first. When there is only one IPsec terminator, regardless of how many hubs are connected to the same IPsec terminator, it is not possible to designate one hub as the primary hub.

For more detailed information on selecting devices for a VPN, see Selecting Devices for Your VPN Topology, page 25-32.

Step 5 Click **Finish** to close the wizard and start the rediscovery process. The Discovery Status window opens and displays the status of the rediscovery and indicates whether the rediscovery of each device has been successful or has failed (see Viewing Policy Discovery Task Status, page 5-22). Error or warning messages are provided to indicate the source of any problems, which may be VPN specific or device specific.

When the rediscovery process completes successfully, and you close the Discovery Status dialog box, the Site-to-Site VPN Manager window opens, displaying summary information for the VPN that was rediscovered.

Creating or Editing VPN Topologies

Security Manager supports three basic types of topologies with which you can create a site-to-site VPN. Use the Create VPN wizard to create a hub-and-spoke, point-to-point, or full mesh VPN topology across multiple device types. For more information about these topologies, see Understanding VPN Topologies, page 25-2.

Tip If you want to create an Extranet point-to-point VPN, read Creating or Editing Extranet VPNs, page 25-66 instead of this topic.

Creating a VPN topology involves specifying the devices and the networks that make up the site-to-site VPN. You define the devices and their roles (such as hub, spoke, peer, key server, group member), the VPN interfaces that are the source and destination endpoints of the VPN tunnel, and the protected networks that will be secured by the tunnel. When you create a VPN topology, you assign to it the IPsec technology (such as Regular IPSec, IPSec/GRE, GRE Dynamic IP, DMVPN, Large Scale DMVPN, Easy VPN, GET VPN) with which a predefined set of policies is associated. See Understanding Mandatory and Optional Policies for Site-to-Site VPNs, page 25-6.

Note When you complete the Create VPN wizard, your topology might be immediately deployable, because Security Manager provides defaults for mandatory policies. However, if you use Security Manager defaults, be sure to verify that the settings will work properly in your network. For more information, see Understanding and Configuring VPN Default Policies, page 25-12.

When you edit a VPN topology, the Edit VPN dialog box contains the same pages as the Create VPN wizard (except for the VPN defaults page), but the pages are laid out in a tabbed format rather than being presented as a wizard. The only exception is for GET VPN topologies, where you can edit only the name and description of the topology (you must edit GET VPN policies to change topology attributes, see Configuring GET VPN, page 29-12). Clicking **OK** on any tab in the dialog box saves your definitions on all the tabs. For all topologies, you must edit mandatory and optional policies originally presented on the VPN defaults page directly.

By editing a VPN topology, you can change its device structure (adding or removing devices), change the VPN interfaces and protected networks defined for a device, or modify the policies that are assigned to the VPN. For example, if your organization frequently opens new sites, you might need to add spokes to an existing hub-and-spoke VPN while applying all policies of the VPN to the new spokes. Or, you might want to increase resiliency by adding a secondary hub to a VPN that has only one hub. While editing a VPN topology, you might also need to modify the policies assigned to it, for example, to change an IKE algorithm to a more secured one, or to change the DES encryption algorithm for a VPN to make it more secure.

Tip After you create a topology, you cannot change the technology used in the VPN. Instead, you must delete the old VPN and create a new one using the desired technology.

To start the Create VPN wizard, or to edit an existing VPN topology:

- To open the Create VPN wizard, in the Site-to-Site VPN Manager Window or the Site-to-Site VPN policy page (Device View), click the **Create VPN Topology (+)** button and select the type of VPN topology you want to create from the options that are displayed—Hub and Spoke, Point to Point, or Full Mesh. Use the Back and Next buttons to move through the pages; when finished, click Finish to create the topology.
- To open the Edit VPN dialog box, select the VPN topology in the Site-to-Site VPN Manager window or the Site-to-Site VPN policy page (Device View) and click the **Edit VPN Topology (pencil)** button.

The pages or tabs that appear and their sequence depend on the type of VPN topology you are creating, as explained in the following table.

Table 25-4 Create/Edit VPN Wizard Pages

Page	Hub and Spoke VPN	Point to Point VPN	Full Mesh VPN
Name and Technology Page. See Defining the Name and IPsec Technology of a VPN Topology, page 25-30.	Step 1	Step 1	Step 1
Device Selection Page. See Selecting Devices for Your VPN Topology, page 25-32.	Step 2	Step 2	Step 2
Endpoints Page. See Defining the Endpoints and Protected Networks, page 25-34. From this page, you can also create several advanced configurations; see the information following the table for further explanation.	Step 3	Step 3	Step 3 (Regular IPSec, IPSec GRE only)
High Availability Page. See Configuring High Availability in Your VPN Topology, page 25-52.	Step 4	—	—
GET VPN Group Encryption Policy Page. See Defining GET VPN Group Encryption, page 25-54.	—	—	Step 3 (GET VPN only.)

Table 25-4 Create/Edit VPN Wizard Pages (continued)

Page	Hub and Spoke VPN	Point to Point VPN	Full Mesh VPN
GET VPN Peers Page. See Defining GET VPN Peers, page 25-60.	—	—	Step 4 (GET VPN only.)
VPN Defaults Page. See Assigning Initial Policies (Defaults) to a New VPN Topology, page 25-62.	Step 5	Step 4	Step 4 (Step 5 for GET VPN.)
Synchronize Keys dialog box. When completing the Create VPN wizard for a GET VPN, you are asked if you want to synchronize keys. Clicking **Yes** initiates the process. See Generating and Synchronizing RSA Keys, page 29-13.	—	—	Step 6 (GET VPN only.)

Either during or after you create a VPN topology, you can also create the following advanced configurations when editing endpoints:

- VRF-Aware IPsec on a hub in a hub-and-spoke topology (see Configuring VRF Aware IPsec Settings, page 25-48).
- A VPNSM or VPNSPA/VSPA on a Catalyst 6500/7600 in a hub-and-spoke, point-to-point, or full mesh VPN topology (see Configuring VPNSM or VPN SPA/VSPA Endpoint Settings, page 25-42).
- A Firewall Services Module together with a VPN Services Module or VPN SPA on a Catalyst 6500/7600 device in a hub-and-spoke, point-to-point, or full mesh VPN topology (see Configuring a Firewall Services Module (FWSM) Interface with VPNSM or VPNSPA/VSPA, page 25-47).

Note You can create a visual representation of your VPN topology with all its elements in the Map view. For more information, see Creating VPN Topologies in Map View, page 35-21.

Related Topics

- Configuring VPN Topologies in Device View, page 25-19
- Understanding IPsec Technologies and Policies, page 25-5
- Using Wizards, page 1-47

Defining the Name and IPsec Technology of a VPN Topology

Note This topic does not apply to Extranet VPNs. For information about configuring the name of an Extranet VPN, see Creating or Editing Extranet VPNs, page 25-66.

Use the Name and Technology page (or tab) of the Create VPN wizard and Edit VPN dialog box to define a name and description for the VPN topology. When creating a new topology, you must select the IPsec technology that will be assigned to it, but you cannot change the technology when editing an existing topology.

For information on opening the Create VPN wizard or Edit VPN dialog box, see Creating or Editing VPN Topologies, page 25-28.

Note If you are editing an existing VPN, the assigned IPsec technology and type is displayed, but you cannot change them. To change the technology or type, you must delete the topology and create a new one.

The following table describes the options you can configure when defining the name and technology.

Table 25-5 Name and Technology Page

Element	Description
Name	A unique name that identifies the VPN topology.
Description	Information about the VPN topology.
IPsec Technology	The IPsec technology used in the VPN topology: • Regular IPsec • IPsec/GRE • DMVPN (Hub and Spoke VPN only) • Easy VPN (Hub and Spoke VPN only) • GET VPN (Full Mesh VPN only) • Regular IPSec VTI
Type	The technology type field appears if you have selected IPsec/GRE or DMVPN as IPsec technology for a hub-and-spoke topology: • IPsec/GRE—Select either **Standard** (for IPsec/GRE) or **Spokes with Dynamic IP** (for GRE Dynamic IP). For more information, see Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5. • DMVPN—Select either **Standard** (for regular DMVPN) or **Large Scale with IPsec Terminator** (for a large scale DMVPN). For more information, see Configuring Large Scale DMVPNs, page 27-16.

Table 25-5 Name and Technology Page (continued)

Element	Description
IKE version	The Internet Key Exchange (IKE) version to allow in IKE negotiations. When configuring regular IPsec VTI topology, you can allow version 1 (**IKEv1**) or version 2 (**IKEv2**). When configuring regular IPsec topology, you can allow version 1 (**IKEv1**), version 2 (**IKEv2**), or both **IKEv1 & IKEv2**. If you select **IKEv1 & IKEv2**, IKEv1 is automatically used by any device that does not support IKEv2. However, if you select IKEv2 only, you must ensure that you do not select any devices that do not support IKEv2 (the wizard does not prevent an invalid selection). You can edit the IKE Proposal and IPsec Proposal policies to change which IKE versions are supported after creating the VPN if you select the wrong option. For information on IKE and how these versions differ, see Overview of IKE and IPsec Configurations, page 26-2. For information on devices that support IKEv2, see Understanding Devices Supported by Each IPsec Technology, page 25-9. **Tip** When using the Create VPN wizard, if you select an option that allows IKEv2, the wizard never creates a valid topology. After completing the wizard, you must manually configure the IKEv2 Authentication policy to complete the configuration.

Related Topics

- Understanding IPsec Technologies and Policies, page 25-5

Selecting Devices for Your VPN Topology

Note This topic does not apply to Extranet VPNs. For information about selecting devices in an Extranet VPN, see Creating or Editing Extranet VPNs, page 25-66.

Use the Device Selection page (or tab) of the Create VPN wizard and Edit VPN dialog box to select the devices to include in the VPN topology. The contents of this page differ depending on whether you are creating or editing a hub-and-spoke, large scale DMVPN, point-to-point, or full mesh VPN topology. Also, you cannot use this page to edit the membership in a GET VPN (instead, see Configuring GET VPN Group Members, page 29-20 and Configuring GET VPN Key Servers, page 29-18 when working with an existing GET VPN).

For information on opening the Create VPN wizard or Edit VPN dialog box, see Creating or Editing VPN Topologies, page 25-28.

In most cases, the devices that are listed in the **Available Devices** list include only those that can be used for the selected VPN topology type, that support the IPsec technology type, and which you are authorized to view. In addition, the available devices depend on the selected IPsec technology—for example, if the IPsec technology is IPsec/GRE, GRE Dynamic IP, or DMVPN, PIX Firewalls and ASA devices are not displayed. The lists are not adjusted to account for the IKE versions you are supporting in the topology. However for regular IPsec VTI topology configuration, when IKEv1 is selected,

ASA 9.7.1 and above single context devices are displayed; for IKEv2, ASA 9.8.1 and above single context devices are displayed. For more information, see the supported platforms described in Understanding Devices Supported by Each IPsec Technology, page 25-9.

Tip When selecting devices, you can select a device group to select all of the eligible devices in the group.

The following list explains how to add or remove devices based on the type of topology:

- **To select devices for a full mesh VPN topology with Regular IPSec or IPSec/GRE technology**, select them in the Available Devices list and click **>>**.
- **To select devices for a full mesh VPN topology that uses the GET VPN technology:**
  - Select the devices that you want to define as key servers and click **>>** next to the **Key Servers** field.

    If you have more than one key server, use the **Up** and **Down** arrow buttons to ensure the primary key server is listed first. Group members register with the first key server in the list. If the first key server cannot be reached, they try to register with the second key server, and so on.
  - Select the devices that you want to define as group members and click **>>** next to the **Group Members** field.
- **To select devices for a hub-and-spoke VPN topology**:
  - Select the devices that you want to define as hubs (or servers in an Easy VPN configuration) and click **>>** next to the **Hubs** list.

    If you have more than one hub, ensure the hubs list is in priority order with the primary hub listed first. To change the order, select a hub and click the **Up** and **Down** arrow buttons until the device is ordered as desired.

    Note You need to select the primary hub only when there are two or more IPsec terminators. When there is only one IPsec terminator, regardless of how many hubs are connected to the same IPsec terminator, it is not possible to designate one hub as the primary hub.
  - Select the devices that you want to define as spokes (or clients in an Easy VPN configuration) and click **>>** next to the **Spokes** list.
  - If you are configuring a **Large Scale DMVPN with IPsec Terminator** topology, you must also select the Catalyst 6500/7600 devices you want to be **IPsec Terminators** in your Large Scale DMVPN configuration. If you select more than one IPsec Terminator, use the **Up** and **Down** arrow buttons to put them in priority order. For more information, see Configuring Large Scale DMVPNs, page 27-16.
- **To select devices for a point-to-point VPN topology:**
  - From the Devices list, select a device to be **Peer One** and click **>>**.
  - Select another device to be **Peer Two** and click **>>**.
- **To remove devices (any topology or technology combination)**, select them from one of the selected devices lists and click **<<** to move them back to the Available Devices list.

  If you are editing an existing VPN topology, you can remove devices from the VPN topology, but you cannot save your changes if your device selections result in an invalid VPN configuration. When removing devices, you should be aware of the following:

- You cannot remove a device if it is the only hub in a hub-and-spoke VPN topology, unless you replace it with a different hub.
- You cannot remove a device that is one of the two devices in a point-to-point VPN topology, unless you replace it with a different device.
- In a VPN topology with multiple hub devices, deleting a hub causes the appropriate tunnels to be removed.
- If some, but not all, spokes in a VPN topology are deleted, the hub side crypto statements change to reflect the removal.
- GET VPNs must have at least one key server and one group member.

Related Topics

- Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11

Defining the Endpoints and Protected Networks

Use the Endpoints page of the Create VPN wizard and Edit VPN dialog box, or the Peers policy, to view the devices in your VPN topology and to define or edit their VPN characteristics and features. You are primarily defining the external or internal VPN interfaces and the protected networks for the devices in the VPN topology. The VPN interfaces are the interfaces that encrypt the data. The protected networks are the networks that are encrypted.

To get to the Endpoints page, do any of the following:

- Open the Create VPN wizard or the Edit VPN dialog box; for the procedure, see Creating or Editing VPN Topologies, page 25-28.
- In the Site-to-Site VPN Manager, select the desired VPN topology (excepting GET VPN topologies) and select the **Peers** policy.

Tips:

- This configuration applies to all IPsec technology types except GET VPN. To configure GET VPN endpoints when creating the VPN, see Defining GET VPN Peers, page 25-60. For existing GET VPNs, configure endpoints using the Key Servers and Group Members policies; see Configuring GET VPN Key Servers, page 29-18 and Configuring GET VPN Group Members, page 29-20.
- The devices listed on this page are selected in the Device Selection Page (see Selecting Devices for Your VPN Topology, page 25-32). You can change the list only when editing the Peers policy, where you can select a device and click the **Delete (trash can)** button to remove it. To add devices, you must edit the VPN topology itself.
- Although you can edit the endpoints for an Extranet VPN using the Peers policy, you should instead edit the endpoints through the Edit Extranet VPN dialog box by editing the VPN topology. The Endpoints page does not appear in the Create Extranet VPN wizard.

The table shows the role each device plays in the VPN (hub, spoke, peer, or IPsec Terminator), the device name, and the VPN interface and protected networks. Initially, the VPN interface and protected network is set to the default interface roles defined in the Security Manager Administrative settings for external and internal interfaces (see VPN Policy Defaults Page, page 11-74). The endpoint configuration might include configurations not shown in this table, but the VPN interface and protected network are the only required settings.

- To change the endpoint configuration for a device, select it and click the **Edit Row** button beneath the table. You can select more than one device to edit at a time, but the devices must serve the same role, and you cannot include Catalyst 6500/7600 devices or VPN service modules when selecting multiple devices. You perform endpoint editing in the Edit Endpoints Dialog Box, whose content differs depending on the selected device type and IPsec technology.

  See the following topics for detailed information about the options you can configure in the Edit Endpoints dialog box:

  - **VPN Interface tab**—To configure the VPN interface and other required interface settings (see Configuring VPN Interface Endpoint Settings, page 25-36). In some cases, you can also configure dial backup (for more information about dial backup, see Configuring Dial Backup, page 25-40).

    For Catalyst 6500/7600 devices, the VPN Interface tab provides settings that enable you to configure a VPN Services Module (VPNSM) or a VPNSPA/VSPA blade on the device (which might be an IPsec Terminator in a large scale DMVPN), and are described in Configuring VPNSM or VPN SPA/VSPA Endpoint Settings, page 25-42.

    For configuring tunnel based VPN, only the VPN Interface tab appears. Use the **Select** button to choose the tunnel interface.

    Easy VPN works by determining the highest and lowest security level interfaces during ASA bootup. VPN client rejects two or more interfaces having same highest security level. In BVI, Easy VPN determines that there are more than two interfaces with same highest security level because of which VPN client is not enabled. In order to overcome this issue, vpnclient secure interface CLI was introduced for all ASA 5506, 5508, and 5512 [x/h/w] devices from ASA 9.9(2) onwards. Thus, to support the CLI in Cisco Security Manager, starting from version 4.17, a new component “VPN Client Interface” is introduced in Hub & Spoke Topology of type (Easy VPN).

  - **Extranet Device Details**—To configure the endpoint settings for the remote (unmanaged) device in an Extranet VPN. The tab appears in the Peers policy only. Instead of editing the information on this tab, the preferred method is to edit the VPN topology and change the settings there. For more information, see Creating or Editing Extranet VPNs, page 25-66.

  - **Hub Interface tab**—If the selected device is a hub in a large scale DMVPN, specify the interface that is connected to the IPsec Terminator. See Configuring Large Scale DMVPNs, page 27-16.

  - **Protected Networks tab**—To define the networks that are encrypted (see Identifying the Protected Networks for Endpoints, page 25-46). The protected network can be an interface role, network/host group object, or in the case of regular IPsec, an ACL policy object.

  - **FWSM tab**—To define the settings that enable you to connect between a Firewall Services Module (FWSM) and an IPsec VPN Services Module (VPNSM) or VPNSPA/VSPA that is already configured on a Catalyst 6500/7600 device. This is possible only in a hub-and-spoke topology where the hub is a Catalyst 6500/7600 device that has these modules installed. For more information, see Configuring a Firewall Services Module (FWSM) Interface with VPNSM or VPNSPA/VSPA, page 25-47.

  - **VRF Aware IPsec tab**—To configure a VRF-Aware IPsec policy on a hub (IPsec Aggregator) in a hub-and-spoke VPN topology. For more information, see Configuring VRF Aware IPsec Settings, page 25-48 and Understanding VRF-Aware IPsec, page 25-14.

  - **Crypto Map tab**—To manually configure the Crypto Map name and Crypto ACL name for each peer, which is supported by Security Manager starting with version 4.7. Crypto Map and Crypto ACL are supported in regular IPSec technology. Therefore, this configuration is applicable only for the topologies with regular IPSec technology. For more information see, Configuring Crypto Map, page 25-50.

- To view the actual interfaces associated with an interface role for each device, select **Matching Interfaces** in the **Show** list beneath the table. If there are no matching interfaces, "No Match" is displayed. The default is to show the interface role policy object names. To create a valid VPN, these roles must match to actual interfaces defined on the device.

Related Topics

- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Configuring VPN Interface Endpoint Settings

Use the VPN Interface tab in the Edit Endpoints dialog box to edit the VPN interfaces defined for devices in the Endpoints table. When defining a primary VPN interface for a router device, you can also configure a backup interface to use as a fallback link for the primary route VPN interface, if its connection link becomes unavailable. You can configure a backup interface on a Cisco IOS security router, that is in a point-to-point or full mesh topology, or that is a spoke in a hub-and-spoke topology, or is a remote client in an Easy VPN topology. For more information, see Configuring Dial Backup, page 25-40.

Tips

- If the device is a hub in a large scale DMVPN, this tab is called **Hub Interface**. Specify the interface that is connected to the IPsec Terminator in the **Hub Interface Toward the IPsec Terminator** field. Enter the name of the interface or interface role, or click **Select** to select it from a list. For more information, see Configuring Large Scale DMVPNs, page 27-16.
- If the device is a Catalyst 6500/7600 device, the VPN Interface tab provides settings that enable you to configure a VPN Services Module (VPNSM) or a VPNSPA/VSPA blade on the device. For a description of the elements that appear on the VPN Interface tab for a Catalyst 6500/7600 device, see Configuring VPNSM or VPN SPA/VSPA Endpoint Settings, page 25-42. The table below assumes the device is not a Catalyst 6500/7600 device.

Navigation Path

On the Endpoints Page of the Create VPN wizard or Edit VPN dialog box, or on the VPN Peers policy, select a device and click **Edit** to open the Edit Endpoints Dialog Box. Select the **VPN Interfaces** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34.

Field Reference

Table 25-6 Edit Endpoints Dialog Box, VPN Interface Tab

Element	Description
Enable the VPN Interface Changes on All Selected Peers	Available if you selected more than one device on the Endpoints page for editing. When selected, applies any changes you make in the VPN interface tab to all the selected devices.
VPN Interface	The VPN interface defined for the selected device. Enter the name of the interface role policy object that defines identifies the interface, or click **Select** to select it from a list or to create a new interface role object. (See Creating Interface Role Objects, page 6-74.) **Note** When manually configuring Crypto Map in devices, you must specify the IP Address of the peer interface and not its name. If the device is an ASA 5505 version 7.2(1) or later, it must have two interfaces defined with different security levels. For more information, see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.
VPN Client Interface	The VPN client interface defined for the selected device. Click **Select** to select it from a list. From Cisco Security Manager 4.17, you can specify the client interface for the Easy VPN. This is applicable for: • ASA 5506 devices and later • BVI interface or any other physical interface and not for the BVI member interfaces • Devices in hub and spoke topology
VPN Client Secure Interface	Beginning from 4.17, Cisco Security Manager supports ASA 9.9(2)s' EzVPN feature support for BVI. This field allows you to define the secured interface. Select the interface to act as the protected network for tunnel establishment. This feature is applicable only for: • EasyVPN topology • spoke interface • ASA 9.9.2 devices onwards
Connection Type	Only available in a hub-and-spoke VPN topology, if the selected device is an ASA or PIX 7.0+ device, and the selected technology is Regular IPsec. Select the type of connection that the hub or spoke will use during an SA negotiation: • **Answer Only**—To configure the hub to only respond to an SA negotiation, but not initiate it. This is the default for hubs. • **Originate Only**—To configure the device to only initiate an SA negotiation, but not respond to one. This is the default for spokes. • **Bidirectional**—To configure the hub or spoke to both initiate and respond to an SA negotiation.

Table 25-6 Edit Endpoints Dialog Box, VPN Interface Tab (continued)

Element	Description
Local Peer IPSec Termination	Unavailable if the selected technology is Easy VPN. Specifies the IP address of the VPN interface of the local router. You can select one of the following options: • **Tunnel Source IP Address**—Use the IP address of the tunnel source. • **VPN Interface IP Address**—Uses the configured IP address on the selected VPN interface. Only one VPN interface can match the interface role. This option is available only if you select **Configure Unique Tunnel Source for each Tunnel** in the GRE Modes policy. **Note** Beginning with version 4.9, Security Manager enables you to select IPv6 addresses. This feature is supported for interfaces that have IPv6 addresses and is applicable for devices running the ASA software version 9.0 or higher. Also, the option of IPv6 address is available only with Regular IPSec technology. • **IP Address**—Explicitly specify the IP address of the VPN interface of the local router. Use this option when the device is behind a NAT boundary to specify the NAT IP Address. Beginning with version 4.9, Security Manager enables you to specify IPv6 addresses. **Note** If you select a tunnel source as the VPN interface, it is likely that the VPN interface has a dynamically assigned IP address. • **IP Address of Another Existing Interface to be Used as Local Address** (unavailable if IPsec technology is DMVPN)—To use the configured IP address on any interface as a local address, not necessarily a VPN interface. Enter the interface in the field provided. You can choose the required interface by clicking **Select**. A dialog box opens that lists all available predefined interface roles, and in which you can create an interface role object.
Tunnel Source	Available only for IPsec/GRE or DMVPN. If you have enabled the setting to use a unique tunnel source per tunnel interface in the GRE Modes > Tunnel Parameters tab, the **Override Unique Tunnel Source per Tunnel Interface** check box is available. Select this option to specify a different tunnel source for the selected device. Specifies the tunnel source address to be used by the GRE or DMVPN tunnel on the spoke side. You can select one of the following options: • **VPN Interface**—Uses the VPN interface as the tunnel source address. • **Interface**—To use any interface as the tunnel source address, not necessarily a VPN interface. Enter the interface name or click **Select** to select an interface role that identifies the interface (you can also create a role from the selection dialog box).

Table 25-6 Edit Endpoints Dialog Box, VPN Interface Tab (continued)

Element	Description
Dial Backup Settings	
Enable Backup	Available if the selected device is an IOS router that is in a point-to-point or full mesh topology, or that is a spoke in a hub-and-spoke topology, or that is a remote client in an Easy VPN topology. Whether to configure a backup interface to use as a fallback link for the primary route VPN interface, if its connection link becomes unavailable. **Tip** Before configuring a backup interface, you must first configure the dialer interface settings on the device. For more information, see Dialer Interfaces on Cisco IOS Routers, page 62-28.
Dialer Interface	The logical interface through which the secondary route traffic is directed when the dialer interface is activated. This can be a Serial, Async, or BRI interface. Enter the name of the interface or interface role object, or click **Select** to select it from a list.
Primary Next Hop IP Address	Available only if the selected technology is Regular IPsec, IPsec/GRE, GRE Dynamic IP, or Easy VPN. The IP address to which the primary interface connects when it is active. This is known as the next hop IP address. If you do not specify the next hop IP address, Security Manager configures a static route using the VPN interface name. The VPN interface must be point-to-point or deployment fails. You can choose the required IP address by clicking **Select**. The Network/Hosts selector opens, in which you can select a network from which the IP address will be allocated.
Tracking IP Address	The IP address of the destination device to which connectivity must be maintained from the primary VPN interface connection. This is the device that is pinged by the Service Assurance agent through the primary route to track connectivity. The backup connection is triggered if connectivity to this device is lost. If you do not specify an IP address, the primary hub VPN interface is used in a hub-and-spoke or Easy VPN topology. In a point-to-point or full mesh VPN topology, the peer VPN interface is used. You can choose the required IP address by clicking **Select**. The Network/Hosts selector opens, in which you can select a network from which the IP address will be allocated.
Advanced button	Available if the selected technology is Regular IPsec, IPsec/GRE, GRE Dynamic IP, or Easy VPN. Click this button to configure additional optional settings using the Dial Backup Settings Dialog Box, page 25-41.

Configuring Dial Backup

You can use dial backup to provide a fallback link for a primary, direct connection when the primary link becomes unavailable. You can configure dial backup on Cisco IOS security routers that participate in a point-to-point, Extranet, or full mesh VPN topology, or that are spokes in a hub-and-spoke topology. You can also configure it on a remote client router running IOS version 12.3(14)T+ in an Easy VPN topology.

Implementation of the dial backup feature is based on the assumption that two static routes exist:

- A primary route through a primary gateway, which has highest priority.
- A secondary route through a secondary gateway, which has lower priority and only appears in the routing table when the primary gateway is down.

Security Manager configures a logical dialer interface on the spoke. The dialer interface is associated with a physical backup interface. When the primary route is down, the dialer interface is activated and traffic is redirected through this backup interface along the secondary route. To ensure that the spoke-hub traffic is encrypted, Security Manager applies a crypto map to the dialer interface. This crypto map is identical to the crypto map on the VPN interface (the primary route interface). In Easy VPN, the backup configuration is attached to the dialer interface.

Depending on the IOS version, Response Time Reporter (RTR) or Service Level Agreement (SLA) IOS technology is used to detect loss of network performance on the primary route. If the assigned IPsec technology is DMVPN, Dialer Watch-List (DWL) is used.

ISDN Basic Rate Interface (BRI) and analog modem interfaces can be configured as backup interfaces to other primary interfaces. In such a case, an ISDN or analog modem connection is made if the primary interface goes down. Should the primary interface and connection go down, the ISDN or analog modem interface immediately dials out to establish a connection so that network services are not lost.

Before You Begin

- Configure the dialer interface settings on the Cisco IOS routers. This requires defining the relationship between the physical BRI and Async interfaces, and the virtual dialer interfaces used when configuring dial backup. For more information, see Dialer Interfaces on Cisco IOS Routers, page 62-28.
- Make sure that the primary route is functioning.
- For Extranet VPNs, you can configure dial backup on the local (managed) device only.

Step 1 For most VPN topologies, you configure dial backup when creating or editing a site-to-site VPN. You can also edit the Peers policy for existing VPN topologies. For Extranet VPNs, you configure dial backup through the Peers policy only.

Do one of the following:

- In the Create VPN wizard, proceed to the Endpoints page (see Creating or Editing VPN Topologies, page 25-28 and Defining the Endpoints and Protected Networks, page 25-34).
- In the Edit VPN dialog box, click the **Endpoints** tab (see Creating or Editing VPN Topologies, page 25-28 and Defining the Endpoints and Protected Networks, page 25-34).
- For Extranet VPNs, or for editing any other VPN topology, select the Peers policy. For general information on editing endpoints, see Defining the Endpoints and Protected Networks, page 25-34.

Step 2 Select the router on which you want to configure dial backup and click the **Edit (pencil)** button. If there is more than one router that will have the same dialer configuration, you can select and edit them all at once.

This action opens the Edit Endpoints dialog box. Select the **VPN Interface** tab if it is not already selected.

Step 3 On the VPN Interface tab, configure the following options related to dial backup. If you are creating a new VPN, you need to configure the other settings (such as VPN interface) as well. For detailed reference information for these options, see Configuring VPN Interface Endpoint Settings, page 25-36.

- **Enable Backup**—Select this option.
- **Dialer Interface**—Specify the physical interface through which the secondary route traffic will be directed when the logical dialer interface is activated.
- **Primary Next Hop IP Address**—If the selected IPsec technology is Regular IPsec, IPsec/GRE, GRE Dynamic IP, or Easy VPN, enter the next hop IP address. If you do not enter the next hop IP address, Security Manager configures a static route using the interface name.
- **Tracking IP Address**—Specify the IP address of the destination device to which connectivity must be maintained from the primary VPN interface connection. This is the device that is pinged through the primary route to track connectivity. The backup connection is triggered if connectivity to this device is lost.

  If you do not specify an IP address, the primary hub VPN interface is used in a hub-and-spoke or Easy VPN topology. In a point-to-point or full mesh VPN topology, the peer VPN interface is used.

Step 4 If the selected IPsec technology is Regular IPsec, IPsec/GRE, GRE Dynamic IP, or Easy VPN, click **Advanced** to configure additional (optional) settings in the Dial Backup Settings dialog box. These settings are explained in Dial Backup Settings Dialog Box, page 25-41. Click **OK** to save your changes.

Step 5 Click **OK** in the Edit Endpoints dialog box.

Dial Backup Settings Dialog Box

Use the Dial Backup Settings dialog box to define optional settings for configuring a dial backup policy for your site-to-site VPN. These settings are available for Regular IPsec, IPsec/GRE, GRE Dynamic IP, or Easy VPN technologies.

Mandatory settings for dial backup are configured in the VPN Interface tab on the Edit Endpoints dialog box. See Configuring VPN Interface Endpoint Settings, page 25-36.

Note You must configure the dialer interface settings before dial backup can work properly. For more information, see Dialer Interfaces on Cisco IOS Routers, page 62-28.

Navigation Path

To open the Dial Backup Settings dialog box, enable dial backup and click **Advanced** on the **VPN Interface** tab of the Edit Endpoints dialog box. For information on opening the Edit Endpoints dialog box, see Defining the Endpoints and Protected Networks, page 25-34.

Related Topics

- Configuring Dial Backup, page 25-40
- Understanding Easy VPN, page 28-1

Field Reference

Table 25-7 Dial Backup Settings Dialog Box

Element	Description
Next Hop Forwarding Backup Next Hop IP Address	If required, enter the next hop IP address of the ISDN BRI or analog modem backup interface (that is, the IP address to which the backup interface will connect when it is active). You can enter an IP address or the name of a network/host object, or click Select to select a network/host object that specifies the IP address. If you do not enter the next hop IP address, Security Manager configures a static route using the interface name.
Tracking Object Settings	
Timeout	The number of milliseconds the Service Assurance Agent operation waits to receive a response from the destination device. The default is 5000 ms.
Frequency	How often Response Time Reporter (RTR) should be used to detect loss of performance on the primary route. The default is every 60 seconds.
Threshold	The rising threshold in milliseconds that generates a reaction event and stores history information for the RTR operation. The default is 5000 ms.

Configuring VPNSM or VPN SPA/VSPA Endpoint Settings

When you select a Catalyst 6500/7600 device in the Endpoints table for editing, the VPN Interface tab of the Edit Endpoints dialog box provides settings for configuring Cisco VPN Services Modules (VPNSM), Cisco VPN Shared Port Adapters (VPN SPAs), and Cisco VPN Service Port Adapters (VSPAs) on the device. You can select more than one Catalyst 6500/7600 device at the same time. Your changes are applied to all the selected devices.

The device can be in a point-to-point or full mesh VPN topology, or a hub or spoke in a hub-and-spoke VPN topology managed by Security Manager (except in an Easy VPN configuration, where the device cannot be a spoke). These settings must also be configured if the selected device is an IPsec Terminator in a large scale DMVPN, although not all settings described below are available. See Configuring Large Scale DMVPNs, page 27-16.

General Notes

- A Catalyst 6500/7600 device can contain from 3 to 13 chassis slots. Due to the design of the blades, you can install one VPNSM or two VPNSPA/VSPA per slot. The location of a VPNSPA/VSPA is identified with a slot and subslot number. Security Manager stores this information in its inventory, so that Security Manager can manage the VPN topologies.
- If you are configuring intra-chassis high availability, you cannot use a VPNSM blade and a VPNSPA/VSPA blade on the same device as primary and failover blades.
- In a remote access VPN, you can configure only one failover unit for each IPsec proposal. See VPNSM/VPN SPA/VSPA Settings Dialog Box, page 33-6.
- If the Catalyst 6500/7600 has a Firewall Services Module (FWSM), you can configure it to work with these modules. For more information, see Configuring a Firewall Services Module (FWSM) Interface with VPNSM or VPNSPA/VSPA, page 25-47.

- If you are configuring a VPNSM or VPNSPA/VSPA with VRF-Aware IPsec on a device, the device cannot belong to a different VPN topology in which VRF-Aware IPsec is not configured. For more information, see Configuring VRF Aware IPsec Settings, page 25-48.
- Create an inside VLAN on the Catalyst 6500/7600 device, or edit an existing port or VLAN configuration. If the device is configured with VRF-Aware IPsec, you must create a forwarding VLAN.

Notes for VPNSMs

- Security Manager supports the configuration of multiple VPNSMs on a Catalyst 6500/7600 device, but only one module (or two if you are configuring intra chassis high availability) can be configured per VPN topology.
- VPNSM configuration requires that its parent Catalyst 6500/7600 device is running Cisco IOS Software release 12.2(18)SXD1 and later.
- You can use only Layer 3 VLANs for VPNSM configuration.

Notes for VPNSPA/VSPAs

- This configuration also applies if you are configuring an IPsec Terminator in a large scale DMVPN configuration. For more information, see Configuring Large Scale DMVPNs, page 27-16.
- The VPN SPA supports the AES encryption algorithm for all key sizes (128-, 192-, and 256-bit), as well as the DES and 3DES encryption algorithms. For more information, see Deciding Which Encryption Algorithm to Use, page 26-6.

  In VRF mode, the **crypto engine slot** *slot/subslot* {**inside** | **outside**} command is deployed on the inside and outside VPN interfaces.
- Make sure that the Catalyst 6500/7600 device is running Cisco IOS Software release 12.2(18)SXE2 or later.
- If you plan to use Crypto Connect Alternate mode (whereby encrypted traffic entering the VPNSM/VPN SPA is passed through and clear text traffic is bypassed), the Catalyst 6500 device must be running Cisco IOS Software version 12.2(33)SXH or later, and the 7600 router must be running 12.2(33)SRA or later.
- In the case of a DMVPN topology in which multiple hubs participate, if one hub is configured with a VPN SPA blade, a tunnel key must not be configured on *any* of the devices, whether they are spokes or hubs. Devices that participate in such a topology must be running Cisco IOS Software version 12.3T and later in order to support tunnels without keys.

Navigation Path

On the Endpoints Page of the Create VPN wizard or Edit VPN dialog box, or on the VPN Peers policy, select a Catalyst 6500/7600 device, then click **Edit** to open the Edit Endpoints Dialog Box. Select the **FWSM** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34.

Field Reference

Table 25-8 Edit Endpoints Dialog Box, VPN Interface Tab's VPNSM/VPN SPA/VSPA Settings

Element	Description
Enable the VPN Interface Changes on All Selected Peers	**Note** Available if you selected more than one Catalyst 6500/7600 device for editing in the Endpoints page. When selected, applies any changes you make in the VPN interface tab to all the selected devices.
VPNSM/VPN SPA/VSPA Settings	• **Use Crypto Connect Alternate**—When selected, only encrypted traffic entering the VPNSM/VPN SPA on the Catalyst 6500/7600 is passed through. Clear text traffic does not go through (bypasses) the adapters. To use this option, the Catalyst 6500 must be running version 12.2(33)SXH or later, and the 7600 router must be running 12.2(33)SRA or later. This mode is recommended as an alternate to Crypto connect mode for enterprise customers who have a need to support large VPN topologies (financial institutions, for example) or need to pass large amounts of data over an encrypted channel (remote disaster recovery or backup over the Internet). • **Inside VLAN**—The VLAN that serves as the inside interface to the service module or adapter. It is also the hub endpoint of the VPN tunnel (unless VRF-Aware IPsec is configured on the device). Enter the name of the VLAN or interface role object, or click **Select** to select it from a list. • **Slot and Subslot**—The number designating the slot location of the VPNSM or VPNSPA/VSPA. If you are configuring a VPNSPA/VSPA, the subslot number is also required. • **Outside VLAN/External port**—The external port or VLAN that connects to the inside VLAN. Enter the name of the VLAN or interface role object, or click **Select** to select it from a list. You must select an interface or interface role that differs from the one selected for the inside VLAN. **Note** If VRF-Aware IPsec is configured on the device, the external port or VLAN must have an IP address.

Table 25-8 Edit Endpoints Dialog Box, VPN Interface Tab's VPNSM/VPN SPA/VSPA Settings

Element	Description
Tunnel Source	**Note** Available only for a hub when the selected technology is IPsec/GRE or DMVPN. Specifies the tunnel source address to be used by the GRE or DMVPN tunnel on the spoke side. You can select one of the following options: • **Override Unique Tunnel Source per Tunnel Interface**—If you have enabled the setting to use a unique tunnel source per tunnel interface in the GRE Modes > Tunnel Parameters tab, this option is available. Select this option to specify a different tunnel source for the selected device. • **Outside VLAN/External Port (When CCA/VRF is Enabled)**—When the **Use Crypto Connect Alternate** check box is selected, this radio button is available. When selected, specifies the outside VLAN/external port as the tunnel source. • **Inside VLAN**—When selected, uses the interface configured for the inside VLAN as the tunnel source. • **Interface**—To use any interface as the tunnel source address, not necessarily a VPN interface, enter the interface name or click **Select** to select an interface role that identifies the interface. You can create new roles from the selection list.
Local Peer IPSec Termination	Define the IPSec termination point of the VPN interface on the local router: • **Inside VLAN**—Use the interface configured as the inside VLAN. • **IP Address**—Use the IP address of the VPN interface on the local router. Enter the IP address. **Note** If you select a tunnel source as the VPN interface, it is likely that the VPN interface has a dynamically assigned IP address.
Enable Failover Blade	Whether to configure a failover VPNSM or VPNSPA/VSPA blade for intra-chassis high availability. **Note** A VPNSM and VPNSPA/VSPA blade cannot be used on the same device as primary and failover blades. Specify the failover blade, as follows: • **Slot**—The slot number that identifies where the VPNSM blade or VPNSPA/VSPA blade is located. • **Subslot**—If you are configuring a VPNSPA/VSPA, select the number of the subslot (0 or 1) on which the failover VPN SPA blade is installed. **Note** If you are configuring a VPNSM, select the blank option.

Identifying the Protected Networks for Endpoints

Use the Protected Networks tab on the Edit Endpoints dialog box to edit the protected networks that are defined on devices in the Endpoints table. (See Defining the Endpoints and Protected Networks, page 25-34.)

You can specify the protected networks as interface roles whose naming patterns match the internal VPN interface of the device, as network/host group objects containing one or more network or host IP addresses, interfaces, or other network objects, or as access control list objects (if Regular IPsec is the assigned technology).

- If you are editing more than one device at a time, select **Enable the Protected Networks Changes on All Selected Peers** to apply any changes you make in the Protected Networks tab to all the selected devices.
- To add a protected network, select it from the Available Protected Networks list and click **>>** to move it to the Selected Protected Networks list. You can use any combination of interface role objects, network/host group objects (listed in the Protected Networks folder), or Access Control List objects to define the protected network for the device. (ACL objects are available only if Regular IPsec is the assigned technology.)

  Beginning with version 4.9, Security Manager supports IPv6 addresses.

  - The Protected Networks folder now supports IPv6 objects.
  - Access Control Lists folder now supports Extended and Unified ACLs.
  - For Interface Roles, if you select an IPv6 enabled interface and click **>>**, a popup window appears with a list of all the IPv6 addresses that are configured. You can select an address from the list and then click **OK** to move the address to the Selected Protected Networks list. To edit an address, select it in the Selected Protected Networks list and click the **Edit Selection** link.
  - For Extranet VPNs, remote backup peer supports IPv6 addresses.

  Note In a hub-and-spoke VPN topology in which Regular IPsec is the assigned technology, when an ACL object is used to define the protected network on a spoke, Security Manager mirrors the spoke's ACL object on the hub to the matching crypto map entry.

  If you do not provide a crypto map entry, then during deployment Security Manager generates the crypto ACL name on the hub device as the ACL object name on the spoke device appended with an "_1". For example, if the ACL object name of a spoke is, say, "spokeACL", Security Manager generates the Crypto ACL name on the hub device as "spokeACL_1". If there are multiple spoke devices with the same ACL object name, Security Manager generates the crypto ACL name on the hub device as "ACLobjectName_spokeDisplayName_1".

  where, "ACLobjectName" is the ACL object name for all the spoke devices in the topology, and, "spokeDisplayName" is the display name of the spoke devices, which is different for each spoke.

Note Cisco Security Manager creates a new ACL for ASA devices, irrespective of topology type, when you execute any of the following:

- Add an extra entry to a protected network.
- Select Enable Spoke to spoke connectivity check box in the **VPN Global Setting > General Settings** tab for an existing hub and spoke topology.
- Add a new peer (as a spoke) to the existing hub and spoke topology.

This new ACL that is generated on-the-fly may disrupt the VPN traffic. Hence, we recommend you to directly make changes using the ACL building block in protected networks.

- To remove a selected protected network, select it and click the **<<** button.
- If the order of the objects matters, you can adjust the priority order of the selected objects using the Move Up, Move Down buttons to position the objects in the selected list as desired. These buttons are not available if order does not matter.
- If an object that you need to define the protected network is not listed, click the **Create (+)** button to add the object; you are prompted to select the type of object you want to add. You can also modify the definition of an existing object by selecting it and clicking the **Edit (pencil)** button. For more information, see the following topics:
  - Understanding Interface Role Objects, page 6-73 and Creating Interface Role Objects, page 6-74.
  - Understanding Networks/Hosts Objects, page 6-80 and Creating Networks/Hosts Objects, page 6-82.
  - Creating Access Control List Objects, page 6-53.

Navigation Path

On the Endpoints Page of the Create VPN wizard or Edit VPN dialog box, or on the Peers policy, select a device and click **Edit** to open the Edit Endpoints Dialog Box. Select the **Protected Networks** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34.

Configuring a Firewall Services Module (FWSM) Interface with VPNSM or VPNSPA/VSPA

Note From 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any enhancements as FWSM is now End of Life.

Security Manager supports the configuration of a Firewall Services Module (FWSM) with an IPsec VPN Services Module (VPNSM) or VPNSPA/VSPA on a Catalyst 6500/7600 device. This feature enables a FWSM to apply firewall policies to untrusted clients, while the VPNSM or VPN SPA/VSPA provides secure access to the internal network.

Use the FWSM tab on the Edit Endpoints dialog box to define the settings that enable you to connect between the FWSM and a VPNSM or VPNSPA/VSPA that is already configured on a Catalyst 6500/7600 device. The FWSM tab is available only in a hub-and-spoke VPN topology when the selected hub is a Catalyst 6500/7600 device.

Tips

- Before you can define the FWSM settings, you must add the hosting Catalyst 6500/7600 device to the Security Manager inventory and discover its FWSM and its policies and security contexts. See Adding Devices from the Network, page 3-12 and Managing Security Contexts, page 59-7.
- If an inside interface is not already created on the Catalyst 6500/7600 device, you must create it (see Creating or Editing VLANs, page 68-26). Then, you must assign the FWSM inside interface (VLAN) to the appropriate security context, or directly to the FWSM blade.

- You also must configure the settings on the VPN Interfaces tab related to IPsec VPN Services Module (VPNSM) or VPNSPA/VSPA. For more information, see Configuring VPNSM or VPN SPA/VSPA Endpoint Settings, page 25-42.

Navigation Path

On the Endpoints Page of the Create VPN wizard or **Edit** VPN dialog box, or on the VPN Peers policy, select a Catalyst 6500/7600 device that contains an FWSM, then click **Edit** to open the Edit Endpoints Dialog Box. Select the **FWSM** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34.

Field Reference

Table 25-9 Edit Endpoints Dialog Box, FWSM Tab

Element	Description
Enable FWSM Settings	Whether you want to configure the connection between the Firewall Services Module (FWSM) and the VPN Services Module (VPNSM) or VPN SPA on the Catalyst 6500/7600 device.
FWSM Inside VLAN	The VLAN that serves as the inside interface to the Firewall Services Module (FWSM). Enter the name of the interface or interface role, or click **Select** to select it from a list or to create a new interface role object.
FWSM Blade	From the list of available blades, select the blade number to which the selected FWSM inside VLAN interface is connected.
Security Context	If the FWSM inside VLAN is part of a security context (that is, the FWSM is running in multiple-context mode), specify the security context name in this field. The name is case-sensitive.

Configuring VRF Aware IPsec Settings

Use the VRF-Aware IPsec tab on the Edit Endpoints dialog box to configure a VRF-Aware IPsec policy on a hub in your hub-and-spoke VPN topology. You can configure VRF-Aware IPsec as a one-box or two-box solution. For more information about VRF-Aware IPsec, see Understanding VRF-Aware IPsec, page 25-14.

Tips

- VRF-Aware IPsec can be configured only on hubs in a hub-and-spoke VPN topology.
- In a VPN topology with two hubs, you must configure VRF-Aware IPsec on both devices.
- You cannot configure VRF-Aware IPsec on a device that belongs to another VPN topology in which VRF-Aware IPsec is not configured.
- You cannot configure VRF-Aware IPsec on hubs that have been configured with high availability. See Configuring High Availability in Your VPN Topology, page 25-52.
- Deployment might fail if the IPsec Aggregator is configured with the same **keyring** CLI command as the existing preshared key (keyring) command, and is not referenced by any other command. In this case, Security Manager does not use the VRF keyring CLI, but generates the keyring with a different name, causing deployment to fail. You must manually remove the preshared key keyring command through the CLI before you can deploy the configuration.

Navigation Path

On the Endpoints Page of the Create VPN wizard or Edit VPN dialog box, or on the VPN Peers policy, select a device that supports VRF-Aware IPsec configuration in a hub-and-spoke topology, and click **Edit** to open the Edit Endpoints Dialog Box. Select the **VRF-Aware IPsec** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34 and Creating or Editing VPN Topologies, page 25-28.

Field Reference

Table 25-10 Edit Endpoints Dialog Box, VRF Aware IPsec Tab

Element	Description
Enable the VRF Settings Changes on All Selected Peers	Available if you selected more than one device for editing in the Endpoints page. When selected, applies any changes you make in the VRF Settings tab to all the selected devices.
Enable VRF Settings	Whether to enable the configuration of VRF settings on the device. **Note** You can remove VRF settings that were defined for the VPN topology by deselect this check box. However, if VRF-Aware IPsec is configured on a Catalyst 6500/7600 device, disabling it requires additional steps, as explained in Enabling and Disabling VRF on Catalyst Switches and 7600 Devices, page 25-17.
VRF Solution	The type of VRF solution you want to configure: • 1-Box (IPsec Aggregator + MPLS PE)—In the one-box solution, one device serves as the Provider Edge (PE) router that does the MPLS tagging of the packets in addition to IPsec encryption and decryption from the Customer Edge (CE) devices. For more information, see VRF-Aware IPsec One-Box Solution, page 25-14. • 2-Box (IPsec Aggregator Only)—In the two-box solution, the PE device does just the MPLS tagging, while the IPsec Aggregator device does the IPsec encryption and decryption from the CEs. For more information, see VRF-Aware IPsec Two-Box Solution, page 25-15.
VRF Name	The name of the VRF routing table on the IPsec Aggregator. The VRF name is case-sensitive.
Route Distinguisher	The unique identifier of the VRF routing table on the IPsec Aggregator. This unique route distinguisher maintains the routing separation for each VPN across the MPLS core to the other PE routers. The identifier can be in either of the following formats: • *IP address:X* (where *X* is in the range 0- 2147483647). • *N:X* (where *N* is in the range 0-65535, and X is in the range 0-2147483647). **Note** You cannot override the RD identifier after deploying the VRF configuration to your device. To modify the RD identifier after deployment, you must manually remove it using the device CLI, and then deploy again.

Table 25-10 Edit Endpoints Dialog Box, VRF Aware IPsec Tab (continued)

Element	Description
Interface Towards Provider Edge (2-Box solution only.)	The VRF forwarding interface on the IPsec Aggregator towards the PE device. If the IPsec Aggregator (hub) is a Catalyst VPN service module, you must specify a VLAN. Enter the name of the interface or interface role object, or click **Select** to select it from a list or to create a new interface role object.
Routing Protocol (2-Box solution only.)	The routing protocol to be used between the IPsec Aggregator and the PE. The options are BGP, EIGRP, OSPF, RIPv2, or Static route. The default is BGP. If the routing protocol used for the secured IGP differs from the routing protocol between the IPsec Aggregator and the PE, select the routing protocol to use for redistributing the routing to the secured IGP. For information about protocols, see Chapter 61, "Managing Routers". **Note** In a one-box solution, these fields are unavailable as you do not need to specify the routing protocol and AS number. In the one-box solution, only the BGP protocol is supported.
AS Number (2-Box solution, BGP or EIGRP routing only.)	The number that will be used to identify the autonomous system (AS) area between the IPsec Aggregator and the PE. The AS number must be within the range 1-65535. If the routing protocol used for the secured IGP differs from the routing protocol between the IPsec Aggregator and the PE, enter an AS number that will be used to identify the secured IGP into which the routing will be redistributed from the IPsec Aggregator and the PE. This is relevant only when IPsec/GRE or DMVPN are applied.
Process Number (2-Box solution, OSPF routing only.)	The routing process ID number that will be used to identify the secured IGP if you are using OSPF routing. The range is 1-65535.
OSPF Area ID (2-Box solution, OSPF routing only.)	The ID number of the area in which the packet belongs. You can enter any number from 0-4294967295. **Note** All OSPF packets are associated with a single area, so all devices must have the same area ID number.
Next Hop IP Address (2-Box solution, static routing only.)	The IP address of the Provider Edge (PE) or the interface that is connected to the IPsec Aggregator, if you are using static routing.
Redistribute Static Route (2-Box solution, non-static routing only.)	Whether to have static routes advertised in the routing protocol configured on the IPsec Aggregator towards the PE device.

Configuring Crypto Map

Beginning with version 4.7, Security Manager enables you to manually configure the Crypto Map name and Crypto ACL name for each peer device in a VPN topology. This feature is supported only in Regular IPsec topologies.

Use the Crypto Map tab in the Edit Endpoints dialog box to list the peer devices along with the Crypto Map Name and Crypto ACL Name configured for the peers. Selecting any peer device in the list and clicking the **Edit (pencil)**button opens the Edit Crypto Map Entry dialog box.

Note If the topology supports dynamic Crypto Map, the dialog box that opens on clicking the **Edit** button enables you to enter the Dynamic Crypto Map Name.

Navigation Path

On the Edit Endpoints Page of the Create VPN wizard or Edit VPN dialog box, select a device and click **Edit** to open the Edit Endpoints Dialog Box. Select the **Crypto Map** tab in the Edit Endpoints dialog box. For information on how to access these pages and dialog boxes, see Defining the Endpoints and Protected Networks, page 25-34.

Field Reference

Table 25-11 Edit Endpoints Dialog Box, Crypto Map Tab

Element	Default Value
Crypto Map Name	There is no default value. If you do not enter any value, Security Manager uses the Crypto Map Name of the device or generates a new Crypto Map Name. If a Crypto Map already exists on the VPN interface, Security Manager will reuse the same name.
Crypto Map Sequence	Security Manager displays the sequence number of the device in this field after it has discovered the device in the managed network. You cannot edit this value. If you are adding a new VPN topology, Security Manager populates the Sequence Number field with a value of #. You cannot edit this value.
Crypto ACL Name	There is no default value. If you do not enter any value, Security Manager generates a new Crypto ACL name.
Dynamic Crypto Map Name	There is no default value. If you do not enter any value, Security Manager uses the Crypto Map Name of the device or generates a new Crypto Map Name.

- You can apply only one crypto map to an interface.
- You cannot assign the same Crypto Map Name on multiple interfaces of a device.
- You cannot assign different Crypto Map names on the same interface of a device.

Edit Crypto Map Entry Dialog Box

Field Reference

Table 25-12 Edit Crypto Map Entry Dialog Box

Element	Default Value
Crypto ACL Name	There is no default value. If you do not enter any value, Security Manager generates a new Crypto ACL name.

Table 25-12 Edit Crypto Map Entry Dialog Box

Element	Default Value
Crypto Map Sequence	Security Manager displays the sequence number of the device in this field after it has discovered the device in the managed network. You cannot edit this value. If you are adding a new VPN topology, Security Manager populates the Sequence Number field with a value of #. You cannot edit this value.
Crypto Mode	Beginning with Security Manager version 4.12 for ASA devices version 9.6(2) or later, you can select an option from the following Crypto Modes: • Tunnel - Default value. Encapsulation mode will be tunnel mode. • Transport - Encapsulation mode will be Transport mode with option to fallback on tunnel mode, if peer does not support it. • Transport-Require - Encapsulation mode will be Transport-Require mode only. **Note** Transport and Transport-Require modes are supported only for IKEv2"

Configuring High Availability in Your VPN Topology

Use the High Availability page of the Create VPN wizard and Edit VPN dialog box to define a group of hubs as a high availability (HA) group. Configuring high availability is optional.

For information on opening the Create VPN wizard or Edit VPN dialog box, see Creating or Editing VPN Topologies, page 25-28.

High Availability (HA) policies provide automatic device backup when configured on Cisco IOS routers or Catalyst 6500/7600 devices that run IP over LANs. You can configure high availability in a hub-and-spoke VPN topology that uses Regular IPsec or Easy VPN technologies.

In Security Manager, HA is supported by an HA group made up of two or more hub devices that use Hot Standby Routing Protocol (HSRP) to provide transparent, automatic device failover. By sharing a virtual IP address, the hubs in the HA group present the appearance of a single virtual device or default gateway to the hosts on a LAN. One hub in the HA group is always active and assumes the virtual IP address, while the others are standby hubs. The hubs in the group watch for hello packets from active and standby devices. If the active device becomes unavailable for any reason, a standby hub takes ownership of the virtual IP address and takes over the hub functionality. This transfer is seamless and transparent to hosts on the LAN and to the peering devices.

Keep the following points in mind when working with HA groups:

- You can configure High Availability only on hubs in a hub-and-spoke VPN topology that uses Regular IPsec or Easy VPN technologies.
- You can configure high availability only on Cisco IOS routers or Catalyst 6500/7600 devices; however, an HA group cannot contain both Cisco IOS routers and Catalyst 6500/7600 devices.
- If you want to configure stateful failover, the HA group can contain only two hubs, and they must be Cisco IOS routers. You cannot use Catalyst 6500/7600 devices.
- You cannot configure High Availability on hubs that have been configured with VRF-Aware IPsec. See Understanding VRF-Aware IPsec, page 25-14.

- You cannot configure GRE on an HA group.
- A device in an HA group can belong to more than one hub-and-spoke topology.
- A device configured as a hub in a site-to-site VPN with an HA configuration cannot be configured as a hub in a different site-to-site VPN with an HA configuration using the same outside interface. Similarly, such a device cannot be configured as a remote access VPN server on which HA is configured using the same outside interface.
- The same auto-generated preshared key must be used for authentication on all peers. If you specified not to use this option when configuring a preshared key policy, this is overridden during configuration of High Availability. For more information, see Configuring IKEv1 Preshared Key Policies, page 26-48.
- During generation of configurations, all hubs in the HA group receive the same commands, which must be deployed to the HA group as a unit. You cannot deploy to individual hubs in the group.

The following table describes the options for configuring high availability.

Table 25-13 High Availability Page

Element	Description
Enable	Whether to enable high availability configuration on a group of hubs. If you already configured high availability, you can remove the configuration by deselecting this option.
Inside Virtual IP	The IP address that is shared by the hubs in the HA group and that represents the inside interface of the HA group. The virtual IP address must be on the same subnet as the inside interfaces of the hubs in the HA group, but must not be identical to the IP address of any of these interfaces. **Note** If there is an existing standby group on the device, make sure that the IP address you provide is different from the virtual IP address already configured on the device.
Inside Mask	The subnet mask for the inside virtual IP address.
VPN Virtual IP	The IP address that is shared by the hubs in the HA group and represents the VPN interface of the HA group. This IP address serves as the hub endpoint of the VPN tunnel. **Note** If there is an existing standby group on the device, make sure that the IP address you provide is different from the virtual IP address already configured on the device.
VPN Mask	The subnet mask for the VPN virtual IP address.
Hello Interval	The duration in seconds (within the range of 1-254) between each hello message sent by a hub to the other hubs in the group to indicate status and priority. The default is 5 seconds.
Hold Time	The duration in seconds (within the range of 2-255) that a standby hub will wait to receive a hello message from the active hub before concluding that the hub is down. The default is 15 seconds.
Standby Group Number (Inside)	The standby number of the inside hub interface that matches the internal virtual IP subnet for the hubs in the HA group. The number must be within the range of 0-255. The default is 1.

Table 25-13 High Availability Page (continued)

Element	Description
Standby Group Number (Outside)	The standby number of the outside hub interface that matches the external virtual IP subnet for the hubs in the HA group. The number must be within the range of 0-255. The default is 2. **Note** The outside standby group number must be different from the inside standby group number.
Enable Stateful Failover	Whether to enable stateful failover, which uses Stateful SwitchOver (SSO) to ensure that state information is shared between the HSRP devices in the HA group. If a device fails, the shared state information enables the standby device to maintain IPsec sessions without having to re-establish the tunnel or renegotiate the security associations. You can configure stateful failover only on an HA group that contains two hubs that are Cisco IOS routers. This check box is disabled if the HA group contains more than two hubs. In an Easy VPN topology, this check box appears selected and disabled, as stateful failover must always be configured. **Tips:** • When deselected in a Regular IPsec topology, stateless failover is configured on the HA group. Stateless failover will also be configured if the HA group contains more than two hubs. You can configure stateless failover on Cisco IOS routers or Catalyst 6500/7600 devices. • Stateful failover cannot be used when RSA Signature is the IKE authentication method. • Stateful failover can be configured together with PKI configuration, but only on devices with Cisco IOS version 12.3(14)T and later.

Related Topics

- Hub-and-Spoke VPN Topologies, page 25-2
- Understanding Easy VPN, page 28-1

Defining GET VPN Group Encryption

Use the GET VPN Group Encryption page to define the group settings and security associations for a GET VPN topology.

The contents of this page differ depending on whether you are using the Create VPN wizard or you are editing the Group Encryption Policy. The wizard page is not tabbed, whereas the policy is tabbed. There is an extra field on the wizard page to allow the security association configuration.

To open the GET VPN Group Encryption page:

- When creating a new GET VPN, use the Create VPN wizard. For information on starting the wizard, see Creating or Editing VPN Topologies, page 25-28.

- (Site-to-Site VPN Manager Window) Select an existing GET VPN topology and then select **Group Encryption Policy** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > Group Encryption Policy**, and then select an existing policy or create a new one.

The following table describes the options you can configure when defining the GET VPN group encryption settings.

Table 25-14 GET VPN Group Encryption Policy Page

Element	Description
Group Settings Tab	
Group Name	The name of the Group Name of Interpretation (GDOI) group. This name is the same as a VPN name.
Group Identity	A parameter that is used to identify the group. All key servers and group members use this parameter to identify with the group. The identity can be either a number (such as 3333) or any IP address (such as the multicast address used for rekey).
Receive Only	If enabled, group members decrypt traffic and forward it in clear text. This feature is useful for testing the VPN. In normal operation, ensure that this option is not selected. For detailed information, see Using Passive Mode to Migrate to GET VPN, page 29-23.
Security Policy (Create VPN wizard only.)	An ACL policy object to be used as the security policy. For a detailed explanation of the contents of this object and how it relates to the group member security policy, see Understanding the GET VPN Security Policy and Security Associations, page 29-10. This field appears only if you are using the Create VPN wizard. In the Group Encryption Policy, you configure the security policy on the Security Associations tab (described below). **Note** If you are using multicast as the method to distribute the keys, then the ACL policy object must contain a deny rule (ACE) for the multicast address. In this way, the rekey packets sent using multicast will not be encrypted by the TEK. This statement allows the group members to receive rekey packets sent using the multicast protocol.

Table 25-14 GET VPN Group Encryption Policy Page (continued)

Element	Description
Authorization Type	The type of authorization mechanism used by the group: None, Certificates, or Preshared Key. Selecting Certificates or Preshared Key provides additional security in allowing only authorized group members to register with the key server. This type of additional security is required when a key server serves multiple GDOI groups. If you select **Certificates**, you must create a list of certificate filters (using some combination of distinguished name or full-qualified domain name attributes). This filter, located on the key server, specifies the attributes and values used to validate whether the group member is authorized to join the GDOI group. Enter a name for the certificate filter, click the **Add Row (+)** button, and fill in the Add Certificate Filter Dialog Box, page 25-58. **Note** To configure certificate authorization, you must also configure a Public Key Infrastructure (PKI) policy for the GET VPN. The PKI enrollment object that you use should define the same distinguished names, or include the device's fully-qualified domain name, as appropriate. If you select **Preshared Key**, also select an ACL policy object to identify the authorized group members. Use permit rules to identify the host or network addresses of group members.
Key Distribution	The transport method to be used to distribute keys to each group member, either unicast or multicast. For help deciding which to use, see Choosing the Rekey Transport Mechanism, page 29-6. If you select **unicast**, the key server sends a rekey message to each registered group member and waits for an acknowledgment. If you select multicast, the key server sends a rekey message to all group members at once and does not wait for acknowledgment. Rekey messages are retransmitted after the retransmit interval configured in this policy. If you select **multicast**, make sure that the router used as the key server is multicast enabled, and also configure the following options: • **Group IP Address**—The IP address of the multicast group to be used for key distribution. • **Use Static IGMP Joins on Group Members**—If you select this option, the static Source Specific Multicast (SSM) mappings are enabled, which reveal the source of multicast traffic to the group member. In the case of GET VPN, the group member learns the key server address.

Table 25-14 GET VPN Group Encryption Policy Page (continued)

Element	Description
RSA Key Label	The label for the RSA key, which is used to encrypt a variety of messages. This key can either already exist on the device, or it can be an unused new label. If you are creating a new VPN, you are asked at the end of the Create VPN wizard whether you want this key synchronized among the key servers; if you click **Yes**, Security Manager generates the key if it does not already exist. If you change this value for an existing GET VPN, you need to synchronize keys from the Key Servers policy. For more detailed information about how this key is used, and the key generation and synchronization process, see Generating and Synchronizing RSA Keys, page 29-13.
Lifetime (KEK)	The number of seconds that the key encryption key (KEK) is valid. This key is used for encrypting rekey messages. Before the end of this lifetime, the key server sends rekey messages to the group, which includes a new KEK encryption key and transforms and new TEK encryption keys and transforms. The KEK lifetime value should be greater than the TEK lifetime value (it is recommended that the KEK lifetime value be at least three times greater than the TEK lifetime value). The default value of 86,400 seconds is usually appropriate. The TEK lifetime value is configured for each security association (see Add New or Edit Security Association Dialog Box, page 25-58).
Encryption Algorithm	The algorithm used to encrypt the rekey message from the key server to the group member.
Retransmits	The number of times the rekey message can be sent if one or more group members do not receive it.
Interval	The number of seconds between retries.
Security Associations Tab	
Security Associations table	Use the Security Associations table to define security associations for the VPN. The columns in the table summarize the settings for an entry and are explained in Add New or Edit Security Association Dialog Box, page 25-58. When creating a new VPN, the Security Policy field (explained above) is used instead of this tab, which does not appear in the wizard. To configure security associations: • Click the **Add** button to add an entry to the table, and fill in the Add New Security Association dialog box. • Select an entry and click the **Edit** button to edit an existing entry. • Select an entry and click the **Delete** button to delete it.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2

- Configuring GET VPN, page 29-12

Add Certificate Filter Dialog Box

Use the Add Certificate Filter dialog box to define a certificate filter for the group encryption policy for GET VPNs. This filter, located on the key server, specifies the attributes and values used to validate whether the group member is authorized to join the group.

Select one of the following filter types:

- **dn**—(Distinguished name.) Specify a comma separated list of *name=value* pairs in the **Subject** field. For example, OU=Cisco, C=US. When you configure the Public Key Infrastructure policy, the PKI enrollment object you select should define the same values on the Certificate Subject Name tab (see PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66). Using a distinguished name can let you match multiple devices with a single filter.
- **fqdn**—(Fully-qualified domain name.) Specify the fully qualified domain name of a single device (for example, router1.example.com) in the **Domain Name** field. When you configure the Public Key Infrastructure policy, the PKI enrollment object you select should have the **Include Device's FQDN** option selected. Because each device has a unique name, an FQDN filter matches a single device only.

Tip To configure certificate authorization, you must also configure a Public Key Infrastructure (PKI) policy for the GET VPN. The PKI policy is configured on all devices in the VPN.

Navigation Path

From the Group Settings tab on the GET VPN Group Encryption page, select Certificates as the authorization type and click the **Add Row** button under the Authorization Filter table, or select a filter and click the **Edit Row** button. For information on opening the Group Encryption page, see Defining GET VPN Group Encryption, page 25-54.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Add New or Edit Security Association Dialog Box

Use the Add New or Edit Security Association dialog boxes to define an IPSec profile (name and transform set only) and security policy used by the selected GET VPN topology.

Navigation Path

To open the Add New Security Association dialog box, from the Security Associations tab on the GET VPN Group Encryption page, click the **Add Row (+)** button or select an existing association and click the **Edit Row (pencil)** button. For information on opening the Group Encryption page, see Defining GET VPN Group Encryption, page 25-54.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2

- Configuring GET VPN, page 29-12

Field Reference

Table 25-15 Add New Security Association Dialog Box

Element	Description
ID	The sequence number of the profile. This number defines the relative priority of the security association (1 being the highest). If you have more than one security association, the ACLs for each are concatenated (and merged) in the order represented by this number, and the group members process the collected ACL as a single ACL. Keep the default number or enter a new one.
IPSec Profile Name	The name of the IPSec profile.
Transform Sets	The transform set policy objects (security protocols, algorithms, and other settings) defined for the IPSec profile. Separate multiple entries with commas, and place them in priority order. Click **Select** to choose from a list of predefined transform sets or to create a new one.
Security Policy	The access control list policy object defined for the security association. Click **Select** to choose from a list of predefined ACL objects or to create a new one. For a detailed explanation of the contents of this object and how it relates to the group member security policy, see Understanding the GET VPN Security Policy and Security Associations, page 29-10. **Note** If you are using multicast as the method to distribute the keys, then the ACL policy object must contain a deny rule (ACE) for the multicast address. In this way, the rekey packets sent using multicast will not be encrypted by the TEK. This statement allows the group members to receive rekey packets sent using the multicast protocol.

Table 25-15 Add New Security Association Dialog Box (continued)

Element	Description
Enable Anti-Replay	Whether to enable the anti-replay feature, which helps prevent eavesdroppers from inserting packets into the data stream. You can configure anti-replay based on traffic counters or time: • **Counter Window Size**—Although this is the default, it is not recommended. Counter-based anti-replay is useful only if there are two group members (essentially a point-to-point VPN). Select a window size. • **Time Window Size**—This is the preferred method, but it requires that there are more than two group members. Enter the number of seconds of the interval duration of the Synchronous Anti-Replay (SAR) clock. The value range is 1 through 100. The default value is 100. For more information on time-based anti-replay, see Understanding Time-Based Anti-Replay, page 29-11. **Note** If you are encrypting high packet rates for count-based anti-replay, ensure that you do not make the KEK or TEK lifetime too long or it can take several hours for the sequence number to wrap. For example, if the packet rate is 100 kilopackets per second, the lifetime should be configured as less than 11.93 hours so that the SA is used before the sequence number wraps.
Enable IPSec Lifetime	Whether to configure an IPsec security association lifetime that overrides the global setting, which is configured in the Global Settings for GET VPN policy (see Configuring Global Settings for GET VPN, page 29-16). This lifetime value controls how long the traffic encryption key (TEK) can be used before a rekey is required. Configure a value based on the volume of traffic (in kilobytes) between group members, seconds, or both. The key expires when either of the values is reached. Use the following recommendations: • The lifetime should be significantly shorter than the one used for the key encryption key (KEK) (see Defining GET VPN Group Encryption, page 25-54), perhaps a third of the length. • The timed lifetime is the recommended approach, because high traffic volumes can cause excessive rekeys (with potential data loss). • Leave a field blank to not override that global setting.

Defining GET VPN Peers

Use the GET VPN Peers page of the Create VPN wizard to configure peer properties for the key servers and group members in a GET VPN topology. After creating the topology, use the **Key Servers** and **Group Members** policies to modify these settings. The policies are the same as the wizard page, except that the key server and group member tables are split into separate policies.

Tip The list of key servers and group members includes those devices you selected on the Device Selection page of the wizard (see Selecting Devices for Your VPN Topology, page 25-32), however, you can use the **Add (+)** and **Delete (trash can)** buttons to add or remove devices from this page.

Examine the list of key servers and group members to determine if the default settings are appropriate for your VPN. You can select **Matching Interfaces** from the **Show** field below each table to display the actual interfaces that will be selected by the default interface roles. The interface roles must resolve to actual interfaces on the device for the GET VPN configuration to be valid.

Before You Begin

This procedure describes how to define peers for GET VPN when creating a new VPN, and explains just the GET VPN peers configuration. For information on opening the Create VPN wizard, see Creating or Editing VPN Topologies, page 25-28.

Related Topics

- Configuring Fail-Close to Protect Registration Failures, page 29-8
- Using Passive Mode to Migrate to GET VPN, page 29-23
- Configuring GET VPN Key Servers, page 29-18
- Configuring GET VPN Group Members, page 29-20

Step 1 Configure the key servers if the default settings are not appropriate.

For each key server you want to modify, select it, click the **Edit (pencil)** button beneath the table, and configure at least following settings. For information on all available settings, see Edit Key Server Dialog Box, page 29-19.

- **Identity Interface**—Select the interface that group members use to identify the key server and register with it. The default is the Loopback interface role, which identifies all Loopback interfaces defined on the key server.
- **Priority**—Define the role of the key server as primary or secondary by entering a priority value between 1-100. The key server with the highest priority becomes the primary key server. If two or more key servers are assigned the same priority value, the device with the highest IP address is used. The default priority is 100 for the first key server, 95 for the second, and so on.

Note There can be more than one primary key server if the network is partitioned.

Step 2 Move key servers up or down in the table to specify the order that group members use to register with key servers. Group members register with the first key server in the list. If the first key server cannot be reached, they will register with the second key server, and so on. Note that this order does not define the overall key server priority, which is used to determine which key server is the primary key server.

Step 3 Configure the group members if the default settings are not appropriate.

For each group member you want to modify, select it, click the **Edit (pencil)** button beneath the table, and configure at least the following settings:

- **GET-Enabled Interface**—This is the VPN-enabled outside interface to the provider edge (PE). Traffic originating or terminating on this interface is evaluated for encryption or decryption, as appropriate. You can configure multiple interfaces by selecting an interface role object that resolves to more than one interface. Click **Select** to select an interface role object or to create a new object.

- **Interface To Be Used As Local Address**—The interface whose IP address is used to identify the group member to the key server for sending data, such as rekey information. If GET is enabled on only one interface, you do not need to specify the interface to be used as the local address. If GET is enabled on more than one interface, you must specify the interface to be used as the local address. Enter the name of the interface or interface role, or click **Select** to select an interface role.

For information on the other available settings, see Edit Group Member Dialog Box, page 29-21.

Assigning Initial Policies (Defaults) to a New VPN Topology

Use the VPN Defaults page of the Create VPN wizard to view and select the shared site-to-site VPN policies that will be assigned to the VPN topology you are creating. The page displays all the available mandatory and optional policies that can be assigned to your VPN topology, according to the selected IPsec technology. (For more information, see Understanding Mandatory and Optional Policies for Site-to-Site VPNs, page 25-6.)

For information on opening the Create VPN wizard, see Creating or Editing VPN Topologies, page 25-28. After you create the topology, you edit these policies directly.

For each policy type, select the shared VPN policy you want to assign to your VPN topology. Only shared policies are available for selection. Use the following tips to guide your selection:

- The initial defaults listed in this page are configured in the Security Manager Administration VPN Policy Defaults Page, page 11-74. If no specific default was configured for a mandatory policy, the Factory Default policy is selected. For more information about configuring default policies, see Understanding and Configuring VPN Default Policies, page 25-12.
- The shared policies listed are only those that have been committed to the database. For example, if you create a new shared IPsec Proposal policy before using the Create VPN wizard, but you do not submit (and have approved, if necessary) the policy beforehand, the new policy does not appear in the list. Ensure that you submit policies before creating a VPN if you need to use the new policies.
- If a policy is mandatory, you must make a selection. If there are no shared policies, Factory Default is your only option. You can always edit the policy after you create the topology.

Note If you try to select a shared policy that is currently locked by another user, a message is displayed warning you of a lock problem. To bypass the lock, select a different policy or cancel the VPN topology creation until the lock is removed. For more information, see Understanding Policy Locking, page 5-8.

- If a policy is optional, and there are no shared policies, you cannot select anything. If you want the features provided by that policy, configure it after you finish creating the topology.
- To view the contents of the policy in a read-only dialog box, select the policy and click the **View Contents** button beside the policy list.
- If you are creating a topology that supports IKEv2 only, the Create VPN wizard will still create either an IKEv1 Preshared Key or IKEv1 Public Key Infrastructure policy according to your selection. There are no default configurations for IKEv2 Authentication policies. Whenever you choose to support IKEv2, you must manually edit the IKEv2 Authentication policy after creating the VPN to define at least the global IKEv2 settings. You can also create peer-specific IKEv2 overrides. When supporting IKEv2 only, you can unassign the IKEv1-specific policies created by the wizard.

When you are done, click **Finish** to create the new VPN topology. The new VPN topology appears in the VPNs selector in the Site-to-Site VPN window, with the VPN Summary page displayed. See Viewing a Summary of a VPN Topology's Configuration, page 25-63.

Viewing a Summary of a VPN Topology's Configuration

Use the VPN Summary page to view a summary of the configuration of a selected VPN topology. This includes information about the type of VPN topology, its devices, the assigned technology, and specific policies that are configured in it. The summary page is opened automatically after you create a VPN topology. When creating an Extranet VPN, it is also shown as the final step of the Create Extranet VPN wizard.

To open the VPN Summary page for a VPN topology:

- (Site-to-Site VPN Manager Window, page 25-18) Select the VPN topology, then select **VPN Summary** from the Policies list.
- (Device view) Select a device that participates in the VPN and select the **Site-to-Site VPN** policy from the Policies list. Select the VPN topology, then click the **Edit VPN Policies** button. This opens the Site-to-Site VPN Manager window with the topology selected, where you can select **VPN Summary** from the Policies list.

The following table describes the information shown on this page.

Note The summary for standard VPNs is significantly different from the summary for Extranet VPNs. This table is divided in two, with the top half explaining summaries for standard VPNs, and the bottom half explaining summaries for Extranet VPNs.

Table 25-16 VPN Summary Page

Element	Description
Summary Information for Standard VPNs	
Name	The name of the VPN topology.
Technology	The IPsec technology assigned to the VPN topology. See Understanding IPsec Technologies and Policies, page 25-5.
Type	The VPN topology type: Hub-and-Spoke, Point-to-Point, or Full Mesh.
Description	A description of the VPN topology.
IPsec Terminator	Available if the VPN topology is large scale DMVPN. The name of the IPsec Terminators used to load balance GRE traffic to the hubs in the large scale DMVPN.
Primary Hub	Available if the VPN topology type is hub-and-spoke. The name of the primary hub in the hub-and-spoke topology.
Failover Hubs	Available if the VPN topology type is hub-and-spoke. The name of any secondary backup hubs that are configured in the hub-and-spoke topology.
Number of Spokes	Available if the VPN topology type is hub-and-spoke. The number of spokes that are included in the hub-and-spoke topology.

Table 25-16 VPN Summary Page (continued)

Element	Description
Peer 1	Available if the VPN topology type is point-to-point. The name of the device that is defined as Peer One in the point-to-point VPN topology.
Peer 2	Available if the VPN topology type is point-to-point. The name of the device that is defined as Peer Two in the point-to-point VPN topology.
Number of Peers	Available if the VPN topology type is full mesh. The number of devices included in the full mesh VPN topology.
IKE Proposal	The security parameters of the IKEv1 proposal configured in the VPN topology. See Configuring an IKE Proposal, page 26-9. **Note** IKEv2 proposals are not displayed in the summary.
Dynamic VTI	Available in an Easy VPN topology. Displays if a dynamic virtual template interface is configured on a device in an Easy VPN topology. See Configuring Dynamic VTI for Easy VPN, page 28-12.
Transform Sets	The IPsec IKEv1 transform sets that specify the authentication and encryption algorithms that will be used to secure the traffic in the VPN tunnel. See Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22. **Note** IPsec IKEv2 transform sets are not displayed in the summary.
Preshared Key	Unavailable if the selected technology is Easy VPN. Specifies whether the shared key to use in the IKEv1 preshared key policy is user defined or auto-generated. See Configuring IKEv1 Preshared Key Policies, page 26-48. **Note** IKEv2 preshared key settings are not displayed in the summary.
Public Key Infrastructure	If an IKEv1 Public Key Infrastructure policy is configured in the VPN topology, specifies the certificate authority (CA) server. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54. **Note** IKEv2 PKI configurations are not displayed in the summary.
Routing Protocol	Available only if the selected technology is IPsec/GRE, GRE Dynamic IP, or DMVPN. The routing protocol and autonomous system (or process ID) number used in the secured IGP for configuring a GRE, GRE Dynamic IP, or DMVPN routing policy. **Note** Security Manager adds a routing protocol to all the devices in the secured IGP on deployment. If you want to maintain this secured IGP, you must create a router platform policy using this routing protocol and autonomous system (or process ID) number. See Understanding the GRE Modes Page, page 27-1.

Table 25-16 VPN Summary Page (continued)

Element	Description
Tunnel Subnet IP	Available only if the selected technology is IPsec/GRE, GRE Dynamic IP, or DMVPN. If a tunnel subnet is defined, displays the inside tunnel interface IP address, including the unique subnet mask. See Understanding the GRE Modes Page, page 27-1.
User Group	Available for an Easy VPN topology. If a User Group policy is configured on a device in the Easy VPN topology, displays the details of the policy. See Configuring a User Group Policy for Easy VPN, page 28-14.
PIX7.0/ASA Tunnel Group	Available for an Easy VPN topology. If a Connection Profile policy is configured on a PIX Firewall version 7.0+ or ASA appliance in the Easy VPN topology, displays the details of the policy. See Connection Profiles Page, page 31-8.
High Availability	Available if the VPN topology type is hub-and-spoke. If a High Availability policy is configured on a device in your hub-and-spoke VPN topology, displays the details of the policy. See Configuring High Availability in Your VPN Topology, page 25-52.
VRF-Aware IPsec	Available if the VPN topology type is hub-and-spoke. If a VRF-Aware IPsec policy is configured on a hub in your hub-and-spoke VPN topology, displays the type of VRF solution (1-Box or 2-Box) and the name of the VRF policy. See Configuring VRF Aware IPsec Settings, page 25-48.
Summary Information for Extranet VPNs	
IKE Phase 1 Proposal section	The parameters for the IKE Phase 1 proposal, which are defined in the IKE Proposal policy object that is assigned to the Extranet. For information about the settings, see the following topics: • Configuring IKEv1 Proposal Policy Objects, page 26-10 • Configuring IKEv2 Proposal Policy Objects, page 26-14
IKE Phase 2 Proposal section	The parameters of the IKE Phase 2 proposal. Most of these parameters are configured in the IPsec transform set policy object assigned to the Extranet. For explanations, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27. The Lifetime attribute parameter is defined in the VPN Global Settings policy, see Configuring VPN Global Settings, page 26-30. The Perfect Forward Secrecy parameter is defined in the IPsec Proposal policy, see Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22.
Authentication section	The preshared key or the PKI enrollment policy object that defines the certificate used to authenticate the connection. When using preshared keys, you can click the **Show/Hide Key** button to toggle between showing and masking the key. If you print the summary or generate a PDF, the key is shown or hidden based on your selection here.

Table 25-16 VPN Summary Page (continued)

Element	Description
Local	The device at the local (managed) end of the Extranet VPN, including the display name, VPN interface name and IP address, and the protected networks.
Remote	The device at the remote (unmanaged) end of the Extranet VPN, including the device name, the IP address of the VPN interface, and the protected networks.
Print button	Click this button to print the summary. The preshared key is shown or hidden based on what is currently displayed in the page. To print the summary, you must have Adobe Acrobat Reader installed. Security Manager generates a PDF of the summary and then prints it using Acrobat's printing capability.
Generate PDF button	Click this button to create a PDF of the summary. The preshared key is shown or hidden based on what is currently displayed in the page. You are prompted for a file name and a location to save the PDF.

Related Topics

- Configuring an IKE Proposal, page 26-9
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22
- Configuring IKEv1 Preshared Key Policies, page 26-48
- Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54
- Configuring GRE Modes for GRE or GRE Dynamic IP VPNs, page 27-6
- Configuring GRE Modes for DMVPN, page 27-12
- Configuring Large Scale DMVPNs, page 27-16
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring a User Group Policy for Easy VPN, page 28-14
- Configuring a Connection Profile Policy for Easy VPN, page 28-13
- Creating or Editing Extranet VPNs, page 25-66

Creating or Editing Extranet VPNs

Security Manager provides a simplified method of creating a regular IPsec point-to-point VPN between a device that you are managing in Security Manager and one that is not managed. This type of VPN is called an *Extranet*.

Typically, an Extranet is a site-to-site VPN connection between your network and the network of a partner or a service provider. However, it can also be a VPN connection within your organization's network, but between devices managed by different groups, or between a Cisco device and a non-Cisco device (which Security Manager cannot manage).

Use the Create Extranet VPN wizard to create this type of point-to-point VPN topology. Creating an Extranet VPN involves specifying the devices, the VPN interfaces that are the source and destination endpoints of the VPN tunnel, and the protected networks that will be secured by the tunnel. You also specify the IKE and IPsec proposals and preshared key or certificates required to complete a secure connection.

When you edit an Extranet VPN topology, the Edit Extranet VPN dialog box contains the same pages as the Create Extranet VPN wizard, with the exception of the IKE proposal page, but the pages are laid out in a tabbed format rather than being presented as a wizard. Clicking **OK** on any tab in the dialog box saves your definitions on all the tabs. For IKE proposals, IPsec proposals, preshared keys, and Public Key Infrastructure certificates, you must edit the policies directly.

Tips

- VPN default policies do not apply to Extranet VPNs. The settings defined on the Security Manager Administration VPN Defaults page are ignored. If you have shared policies that you want to use in the Extranet VPN configuration, you can assign them to the VPN after you create it with the Create Extranet VPN wizard. Assigning the shared policy replaces the policy created by the wizard.
- You cannot select your pre-defined IKE proposal or IPsec transform set policy objects when creating an Extranet VPN. If you have existing objects that you want to use, you can edit the relevant policies after creating the VPN and select the objects. You can then delete the objects created by the Create Extranet VPN wizard, if desired.
- After creating an Extranet VPN, you cannot convert it to a standard point-to-point VPN, where you are managing both ends of the VPN in Security Manager. Instead, you must delete and recreate the VPN.
- You can configure Extranet VPN connections for regular IPsec point-to-point connections only. For example, you cannot use this method to identify a GET VPN key server that exists in your service provider's network. To configure all other types of Extranet connections, you must add dummy unmanaged devices to the Security Manager inventory as described in Including Unmanaged or Non-Cisco Devices in a VPN, page 25-11.

Related Topics

- Understanding VPN Topologies, page 25-2
- Configuring VPN Topologies in Device View, page 25-19
- Understanding IPsec Technologies and Policies, page 25-5
- Using Wizards, page 1-47

Step 1 Do one of the following

- To create a new Extranet VPN, in the Site-to-Site VPN Manager Window or the Site-to-Site VPN policy page (Device View), click the **Create VPN Topology (+)** button and select **Extranet VPN**. The Create Extranet VPN wizard starts with the Name and Technology page.
- To edit an existing Extranet VPN, select the VPN topology in the Site-to-Site VPN Manager window or the Site-to-Site VPN policy page (Device View) and click the **Edit VPN Topology (pencil)** button. The Edit Extranet VPN dialog box opens to the Device Selection tab.

Step 2 On the Name and Technology page or tab, configure the following; only the name is required:

- **Name**—A unique name that identifies the VPN topology.
- **Description**—A description of the VPN, up to 1024 characters.

- **Creation Date**—The date on which the VPN was created. When creating the VPN, today's date is the default. However, you can click the calendar icon beside the edit box and select the desired date.
- **Ticket Number**—If you use a ticket system, and the action you are taking relates to a tracked requirement, enter the number in this field. Security Manager does not use this number; it is for your internal tracking purposes only.
- **Last Modified By**—The name, user ID, email address, or other indicator of the person who last changed the settings for the VPN. Security Manager does not use this field; it is for your internal tracking purposes only.

In the wizard, click **Next**; in the Edit Extranet VPN dialog box, click the **Device Selection** tab.

Step 3 On the Device Selection page or tab, configure the devices, interfaces, and protected networks for each end of the connection:

- **Local**—This is the device in your managed network. The device must be in the Security Manager inventory. Configure all of these properties:
  - **Device**—Enter the display name of the device or click **Select** to select it from the list of devices in the inventory. You can select ASA 5500 series devices, PIX firewalls, or Cisco IOS routers (including ASRs).
  - **VPN Tunnel Interface**—The name of the interface or interface role that identifies the external interface for the VPN connection. Click **Select** to select an existing interface or interface role, or to create a new interface role.

    When you select an interface or role, the IP address for the matching interface are listed in the drop-down list next to the IP Address field. Beginning with version 4.9, Security Manager supports IPv6 addresses in Extranet VPN. You can see a list of IPv4 and IPv6 addresses, by default the IPv4 address will be displayed. If no address appears, Security Manager could not determine the IP address. Check your configuration or object selection.
  - **Protected Networks**—The networks that the device is protecting for this VPN. Click **Select** to display the Protected Network Selection dialog box in which you can specify the protected networks using an interface name, interface role object, network/host group object, or ACL object. You can also use the Protected Network Selection dialog box to define new network/host group or ACL objects.

    Note You can also edit the local device endpoint settings as described in Defining the Endpoints and Protected Networks, page 25-34. The settings are similar to these, with the added ability to define interface role objects.

  - **Crypto Map name**—You can manually enter the Crypto Map name for the device. There is no default value. If you do not enter any value, Security Manager uses the Crypto Map Name of the device or generates a new Crypto Map Name. If a Crypto Map already exists on the VPN interface, Security Manager will reuse the same name.
  - **Crypto ACL name**—You can manually enter the Crypto ACL name for the device. There is no default value. If you do not enter any value, Security Manager generates a new Crypto ACL name.
  - **Crypto Map Sequence**—Security Manager displays the sequence number of the device in this field after it has discovered the device in the managed network. You cannot edit this value. If you are adding a new VPN topology, Security Manager populates the Sequence Number field with a value of #. You cannot edit this value.

    For more information see, Configuring Crypto Map, page 25-50

- **Crypto Mode**—Beginning with Security Manager version 4.12 for ASA devices version 9.6(2) or later, you can select an option from the following Crypto Modes:
  - Tunnel—Default value. Encapsulation mode will be tunnel mode.
  - Transport—Encapsulation mode will be Crypto mode with option to fallback on tunnel mode, if peer does not support it.
  - Transport-Require—Encapsulation mode will be Transport mode only. Transport Mode is supported only for IKEv2.
- **Remote**—This is the device that you are not managing in Security Manager. Configure all of these properties:
  - **Name**—The name of the device, equivalent to the display name used in the Security Manager inventory.
  - **IP Address**—The IP address of the VPN interface on the device. You can enter a maximum of 10 IP addresses using space as the delimiter. Beginning with version 4.9, Security Manager supports IPv6 addresses in addition to IPv4 addresses.

    Note Beginning from version 4.8, Security Manager enables you to configure multiple peer IP addresses for the same extranet VPN configuration. This allows the next peer device in the list to act as a failover when the first device is not available for VPN services. This backup peer support is available for Cisco Adaptive Security Appliance (ASA) devices and Cisco IOS routers.

  - **Protected Networks**—The networks that the device is protecting for this VPN. Click **Select** to display the Protected Network Selection dialog box in which you can specify the protected networks using a network/host group object or ACL object. You can also use the Protected Network Selection dialog box to define new network/host group or ACL objects.

    Note You can also edit the remote device endpoint settings as described in Defining the Endpoints and Protected Networks, page 25-34. However, the settings are identical to these, and you cannot specify the protected networks using an interface name or interface role object.

In the wizard, click **Next**. In the Edit VPN dialog box, you are finished; to edit the remaining characteristics, you must edit the IKE Proposal, IPsec Proposal, IKEv1 Preshared Key, IKEv1 Public Key Infrastructure, IKEv2 Authentication, and VPN Global Settings policies to change the settings described in the next step.

Step 4 On the IKE Proposal page of the Create Extranet VPN wizard, define the IKE proposal, the IPsec proposal, and either the preshared key or the certificate:

- Select **IKEv1** or **IKEv2**. You can use IKEv2 on ASA 5500 series devices running release 8.4(x) only.

  If you want to change IKE versions after creating the Extranet VPN, you must edit all of these policies to unassign or replace the old configuration while configuring options for the desired version: IKE Proposal, IPsec Proposal, IKEv1 Preshared Key, IKEv1 Public Key Infrastructure, IKEv2 Authentication, VPN Global Settings. For information on how IKEv1 and IKEv2 differ, see Comparing IKE Version 1 and 2, page 26-4.

- Configure the IKE Phase 1 Proposal parameters. These parameters will be used to create an IKE proposal policy object with the name *ExtranetName*_**ikeBB**. For an explanation of the parameters, see Configuring IKEv1 Proposal Policy Objects, page 26-10 or Configuring IKEv2 Proposal Policy Objects, page 26-14.

  To edit these values after creating the VPN, you simply need to edit the object. You can edit the object in the Policy Object Manager or directly through the IKE Proposal policy for the VPN.

  Note The **DH Group** attribute (for Diffie-Hellman modulus group) is called **Modulus Group** in other policies and policy objects.

- Configure the IKE Phase 2 (IPsec) Proposal parameters. Most of these parameters will be used to create an IPsec transform set policy object with the name *ExtranetName*_**transformSet**. For an explanation of the parameters, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27. Note that the AH Hash Algorithm setting is available only if the local device is a router.

  To edit these values after creating the VPN, you simply need to edit the object. You can edit the object in the Policy Object Manager or directly through the IPsec Proposal policy for the VPN.

  The following settings are not part of the IPsec transform set object:

  – **Enable Perfect Forward Secrecy, DH Group**—Whether to use a unique session key for each encrypted exchange, which prevents an attacker from decrypting a captured exchange even if the attacker knows the preshared or private keys used by both ends of the tunnel. If you select this option, also select the Diffie-Hellman (DH) modulus group to use for deriving the key. For more information on the modulus group, see Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7.

    To change this option after creating the VPN, edit the IPsec Proposal policy.

  – **Lifetime**—The number of seconds a security association will exist before expiring. The default is 3,600 seconds (one hour).

    To change this option after creating the VPN, edit the VPN Global Settings policy.

- If you select **Preshared Key** for authentication, enter the key used to authenticate connections with the remote host.

  To edit the key after creating the VPN, you must edit either the IKEv1 Preshared Key or IKEv2 Authentication policy depending on the IKE version you are using. The key is masked in these policies, but you can display the key by selecting the VPN Summary policy and clicking the Show Key button beside the preshared key.

- If you select **Certificate**, select the PKI enrollment object that defines the certificate name. If the required object is not yet defined, select **<Add New>** to open the Add PKI selector, from which you can add new, or edit existing, PKI enrollment objects. For more information about PKI enrollment objects, see PKI Enrollment Dialog Box, page 26-58.

  To edit the certificate settings after creating the VPN, you can edit the object in the Policy Object Manager or directly through either the IKEv1 Public Key Infrastructure or IKEv2 Authentication policy depending on the IKE version you are using.

In the wizard, click **Next**.

Step 5 (Create Extranet VPN wizard only.) On the Summary page, verify that the settings are correct and click **Finish**.

Security Manager creates the topology and the required policy objects, and adds the VPN to the list of VPNs in the Site-to-Site VPN Manager.

Step 6 If you want to configure dial backup, select the **Peers** policy and follow the instructions in Configuring Dial Backup, page 25-40.

Deleting a VPN Topology

Deleting a VPN topology removes IPsec tunnels between peers and all configurations associated with the VPN topology from the devices and networks assigned to the site-to-site VPN. The actual VPN is not removed from the network until you deploy configurations.

Step 1 Do one of the following:

- Select **Manage > Site-To-Site VPNs** to open the Site-to-Site VPN Manager Window, page 25-18.
- In Device view, select a device that participates in the VPN you want to delete, then select the **Site to Site VPN** policy from the policy selector (see Configuring VPN Topologies in Device View, page 25-19).

Step 2 Select the VPN topology you want to delete and click the **Delete VPN Topology (trash can)** button. You are asked to confirm the deletion.

Configuring IKE and IPsec Policies

This chapter describes how to configure Internet Protocol Security (IPsec) and the Internet Security Association and Key Management Protocol (ISAKMP, or IKE) standards to build site-to-site and remote access IPsec Virtual Private Networks (VPNs). These policies are used in regular IPsec and other types of IPsec-based VPN technologies to build VPN tunnels.

Tunneling makes it possible to use a public TCP/IP network, such as the Internet, to create secure connections between remote users and a private corporate network. Each secure connection is called a tunnel.

IPsec-based VPN technologies use the ISAKMP and IPsec tunneling standards to build and manage tunnels. ISAKMP and IPsec accomplish the following:

- Negotiate tunnel parameters.
- Establish tunnels.
- Authenticate users and data.
- Manage security keys.
- Encrypt and decrypt data.
- Manage data transfer across the tunnel.
- Manage data transfer inbound and outbound as a tunnel endpoint or router.

A device in a VPN functions as a bidirectional tunnel endpoint. It can receive plain packets from the private network, encapsulate them, create a tunnel, and send them to the other end of the tunnel where they are unencapsulated and sent to their final destination. It can also receive encapsulated packets from the public network, unencapsulate them, and send them to their final destination on the private network.

The following topics explain the basic IKE and IPsec policies and how to configure them:

- Overview of IKE and IPsec Configurations, page 26-2
- Understanding IKE, page 26-5
- Understanding IPsec Proposals, page 26-18
- Configuring VPN Global Settings, page 26-30
- Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47
- Understanding Public Key Infrastructure Policies, page 26-51
- Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68

Overview of IKE and IPsec Configurations

Internet Key Exchange (IKE) is a key management protocol that is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec. Both phases use proposals when they negotiate a connection.

An IKE proposal is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters will be used to protect subsequent IKE negotiations. For IKE version 1 (IKEv1), IKE proposals contain a single set of algorithms and a modulus group. You can create multiple, prioritized policies at each peer to ensure that at least one policy matches a remote peer's policy. Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation, potentially making it possible to create a single IKE proposal (although you might want different proposals to give higher priority to your most desired options). You can define several IKE proposals per VPN.

You must configure several policies to define the settings required to make successful regular IPsec connections in a site-to-site or remote access VPN. The following procedure provides an overview of the steps required to complete the configuration, and points to other topics that provide detailed information for each step.

Related Topics

- Understanding IKE, page 26-5
- Understanding IPsec Proposals, page 26-18
- Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47
- Understanding Public Key Infrastructure Policies, page 26-51

Step 1 Configure the **IKE Proposal** policy.

In the IKE Proposal policy, you define the IKE proposal policy objects to use for making VPN connections. When defining the IKE proposal object, you select the algorithms to use for encrypting the IKE negotiation and for integrity checking, and the Diffie-Hellman group to use to operate the encryption algorithm. For IKEv1, you also determine whether you are using preshared keys or Public Key Infrastructure, whereas in IKEv2, the IKE proposal does not include a specification for authentication mode.

The following topics explain how to configure the IKE Proposal policy:

- Configuring an IKE Proposal, page 26-9
  - Configuring IKEv1 Proposal Policy Objects, page 26-10
  - Configuring IKEv2 Proposal Policy Objects, page 26-14
- Configuring the IKE Proposal for GET VPN, page 29-15

Step 2 Complete the authentication mode configuration.

Your selection for authentication mode in the IKEv1 proposal, and your decision on which mode to use for IKEv2, controls what other policies are required to complete the authentication mode configuration:

- Preshared keys—For remote access IKEv1 IPsec VPNs, you define the preshared keys in the **Connection Profiles** policy; preshared keys are not supported for IKEv2 in remote access VPNs. For site-to-site VPNs, you define the keys in the **IKEv1 Preshared Keys** or the **IKEv2 Authentication** policy based on the IKE version you are using.

  The following topics explain preshared key configuration:

  - IPSec Tab (Connection Profiles), page 31-19
  - Configuring IKEv1 Preshared Key Policies, page 26-48
  - Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68

- Public Key Infrastructure Certificate Authority servers—If you configure IKE to use Certificate Authority (CA) servers, you must configure the **Public Key Infrastructure** policy. You also use this policy to define the Public Key Infrastructure for SSL VPNs. For site-to-site VPNs, the policy is **IKEv1 Public Key Infrastructure** or **IKEv2 Authentication**, based on the IKE version you are using.

  The Public Key Infrastructure policy identifies the PKI enrollment object that identifies the Certificate Authority server. For site-to-site VPNs, you can select a single PKI enrollment object; for remote access VPNs, you can select all objects needed for your remote access connections. These trustpoints are then identified in the remote access **Connection Profiles** policy (on the IPsec tab).

  The following topics explain public key infrastructure configuration:

  - Understanding Public Key Infrastructure Policies, page 26-51
  - Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54
  - Defining Multiple IKEv1 CA Servers for Site-to-Site VPNs, page 26-55
  - Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56
  - IPSec Tab (Connection Profiles), page 31-19
  - Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68

Step 3 Configure the **IPsec Proposal** policy. The IPsec Proposal policy defines the IPsec transform set policy objects used to create a secure IPsec tunnel for the VPN.

The following topics explain how to configure the IPsec Proposal policy:

- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22
  - Selecting the IKE Version for Devices in Site-to-Site VPNs, page 26-26
  - Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40
- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3

Step 4 Configure the **Global Settings** policy.

The **Global Settings** (remote access) and **VPN Global Settings** (site-to-site) policies define various ISAKMP, IKEv1, IKEv2, IPsec, NAT, fragmentation, and other settings. These settings have default values that are frequently adequate, so normally you need to configure the Global Settings policy only if you want non-default behavior. However, you must configure the policy for remote access IKEv2 IPsec VPNs, because you must specify a remote access global trustpoint on the **IKEv2 Settings** tab.

The following topics explain how to configure the Global Settings policy:

- Configuring VPN Global Settings, page 26-30

  - Configuring VPN Global ISAKMP/IPsec Settings, page 26-33
  - Configuring VPN Global IKEv2 Settings, page 26-37
  - Configuring VPN Global NAT Settings, page 26-42
  - Configuring VPN Global General Settings, page 26-44
- Configuring Global Settings for GET VPN, page 29-16

Step 5 If you are configuring a remote access IKEv2 IPsec VPN, you must also configure several policies for SSL VPN. IKEv2 shares several configuration settings with SSL VPNs. For information on the other policies you need to configure, see Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25.

Comparing IKE Version 1 and 2

There are two versions of IKE: version 1 (IKEv1) and version 2 (IKEv2). When you configure IKE on a device that supports IKEv2, you have the option of configuring either version alone, or both versions together. When the device attempts to negotiate a connection with another peer, it uses whichever versions you allow and that the other peer accepts. If you allow both versions, the device automatically falls back to the other version if negotiations are unsuccessful with the initially chosen version (IKEv2 is always tried first if it is configured). Both peers must support IKEv2 to use it in a negotiation.

Tip Security Manager supports IKEv2 on ASA 8.4(1)+ only. For remote access IPsec VPNs, users must use the AnyConnect 3.0+ client to complete IKEv2 connections, and IKEv2 connections use the same license pool that is used for SSL VPN connections. The legacy VPN Client is used for IKEv1 remote access connections on ASAs. For more information about device support in VPNs, see Understanding Devices Supported by Each IPsec Technology, page 25-9.

IKEv2 differs from IKEv1 in the following ways:

- IKEv2 fixes the Photuris style cookie mechanism.
- IKEv2 has fewer round trips in a negotiation than IKEv1, two round trips versus five for IKEv1 for a basic exchange.
- Transform options are OR'ed, which means that you can specify multiple options in a single proposal rather than creating separate unique proposals for each allowed combination.
- Built-in dead peer detection (DPD).
- Built-in configuration payload and user authentication mode.
- Built-in NAT traversal (NAT-T). IKEv2 uses ports 500 and 4500 for NAT-T.
- Improved re-keying and collision handling.
- A single security association (SA) can protect multiple subnets, which improves scalability.
- Asymmetric authentication in site-to-site VPNs, where each side of a tunnel can have different preshared keys, different certificates, or one side a key and the other side a certificate.
- For remote access IPsec VPNs, you can configure double authentication for IKEv2 connections in the same way that you configure them for remote access SSL VPNs. IKEv1 does not support double authentication.

Related Topics

- Overview of IKE and IPsec Configurations, page 26-2
- Configuring an IKE Proposal, page 26-9

Understanding IKE

Internet Key Exchange (IKE), also called Internet Security Association and Key Management Protocol (ISAKMP), is the negotiation protocol that lets two hosts agree on how to build an IPsec security association (SA). It provides a common framework for agreeing on the format of SA attributes. This includes negotiating with the peer about the SA, and modifying or deleting the SA. IKE creates the cryptographic keys used to authenticate IPsec peers, negotiates and distributes IPsec encryption keys, and automatically establishes IPsec security associations.

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, creating the first tunnel that enables the peers to communicate securely in Phase 2, protecting later ISAKMP negotiation messages. During Phase 2 negotiation, IKE establishes SAs for other applications, such as IPsec, which protects the data sent between peers. Both phases use proposals when they negotiate a connection.

An IKE proposal is a set of algorithms that two peers use to secure the IKE negotiation between them. IKE negotiation begins by each peer agreeing on a common (shared) IKE policy. This policy states which security parameters will be used to protect subsequent IKE negotiations. In remote access IPsec VPNs, you can define several IKE proposals per VPN to create multiple, prioritized policies at each peer to ensure that at least one policy matches a remote peer's policy. For site-to-site VPNs, you can create a single IKE proposal.

To define an IKE proposal, you must specify:

- A unique priority (1 through 65,543, with 1 the highest priority).
- An encryption method for the IKE negotiation, to protect the data and ensure privacy. See Deciding Which Encryption Algorithm to Use, page 26-6.
- A Hashed Message Authentication Codes (HMAC) method (called *integrity algorithm* in IKEv2) to ensure the identity of the sender, and to ensure that the message has not been modified in transit. See Deciding Which Hash Algorithm to Use, page 26-6.
- For IKEv2, a separate pseudo-random function (PRF) used as the algorithm to derive keying material and hashing operations required for the IKEv2 tunnel encryption. The options are the same as those used for the hash algorithm; see Deciding Which Hash Algorithm to Use, page 26-6.
- A Diffie-Hellman group to determine the strength of the encryption-key-determination algorithm. The device uses this algorithm to derive the encryption and hash keys. See Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7.
- An authentication method, to ensure the identity of the peers. See Deciding Which Authentication Method to Use, page 26-8.
- A limit to the time the device uses an encryption key before replacing it.

Tip (ASA devices only.) With IKEv1 policies, for each parameter, you set one value. For IKEv2, you can configure multiple encryption, integrity, PRF, and Diffie-Hellman options. The ASA orders the settings from the most secure to the least secure and negotiates with the peer using that order. This allows you to potentially send a single proposal to convey all the allowed transforms instead of the need to send each allowed combination as with IKEv1.

When IKE negotiation begins, the peer that initiates the negotiation sends all of its policies to the remote peer, and the remote peer searches for a match with its own policies, in priority order.

A match between IKE policies exists if they have the same encryption, hash (integrity and PRF for IKEv2), authentication, and Diffie-Hellman values, and an SA lifetime less than or equal to the lifetime in the policy sent. If the lifetimes are not identical, the shorter lifetime—from the remote peer policy—applies. If no match exists, IKE refuses negotiation and the IKE SA is not established.

The following topics explain how to configure IKE proposals:

- Configuring an IKE Proposal, page 26-9
- Configuring IKEv1 Proposal Policy Objects, page 26-10
- Configuring IKEv2 Proposal Policy Objects, page 26-14
- Configuring the IKE Proposal for GET VPN, page 29-15

Deciding Which Encryption Algorithm to Use

When deciding which encryption and hash algorithms to use for the IKE proposal, your choice is limited to algorithms that are supported by the devices in the VPN.

You can choose from the following encryption algorithms:

- DES (Data Encryption Standard) is a symmetric secret-key block algorithm. It is faster than 3DES and uses less system resources, but it is also less secure. If you do not need strong data confidentiality, and if system resources or speed is a concern, you should choose DES.
- 3DES (Triple DES) is more secure because it processes each block of data three times, each time with a different key. However, it uses more system resources and is slower than DES. 3DES is the recommended encryption algorithm, assuming that the devices support it.
- AES (Advanced Encryption Standard) provides greater security than DES and is computationally more efficient than 3DES. AES offers three different key strengths: 128-, 192- and 256-bit keys. A longer key provides higher security but a reduction in performance. When you configure IKE on a router, the router must use Cisco IOS Software 12.3T or later to use AES.

Note AES cannot be used in conjunction with a hardware encryption card.

Related Topics

- Understanding IKE, page 26-5
- Configuring an IKE Proposal, page 26-9

Deciding Which Hash Algorithm to Use

You can choose from the following hash algorithms. In IKEv2, the hash algorithm is separated into two options, one for the integrity algorithm, and one for the pseudo-random function (PRF).

- SHA (Secure Hash Algorithm) is more resistant to brute-force attacks than MD5. However, it is also more resource intensive than MD5. For implementations that require the highest level of security, use the SHA hash algorithm.

  Standard SHA produces a 160-bit digest.

The following options, which are even more secure, are available for IKEv2 configurations on ASA 8.4(2+) devices:

  - SHA512—A 512-bit key.
  - SHA384—A 384-bit key.
  - SHA256—A 256-bit key.

- MD5 (Message Digest 5) produces a 128-bit digest and uses less processing time for an overall faster performance than SHA, but it is considered to be weaker than SHA.

Related Topics

- Understanding IKE, page 26-5
- Configuring an IKE Proposal, page 26-9

Deciding Which Diffie-Hellman Modulus Group to Use

Security Manager supports the following Diffie-Hellman key derivation algorithms to generate IPsec security association (SA) keys. Each group has a different size modulus. A larger modulus provides higher security, but requires more processing time. You must have a matching modulus group on both peers.

Tip If you select AES encryption, to support the large key sizes required by AES, ISAKMP negotiation should use Diffie-Hellman (DH) Group 5 or higher. For IKEv1, ASA devices support groups 2 and 5 only.

- Diffie-Hellman Group 1: 768-bit modulus. Use to generate IPsec SA keys where the prime and generator numbers are 768 bits.

Note DH group 1 is deprecated and will be removed in later ASA version. In later ASA versions, the default value will be Group 2.

- Diffie-Hellman Group 2: 1024-bit modulus. Use to generate IPsec SA keys where the prime and generator numbers are 1024 bits. Cisco VPN Client Version 3.x or higher requires a minimum of Group 2.
- Diffie-Hellman Group 5: 1536-bit modulus. Use to generate IPsec SA keys where the prime and generator numbers are 2048 bits. Considered good protection for 128-bit keys, but group 14 is better.
- Diffie-Hellman Group 7: Use to generate IPsec SA keys when the elliptical curve field size is 163 characters. Group 7 is not supported on a Catalyst 6500/7600 device with VPNSM or VPN SPA configuration.
- Diffie-Hellman Group 14: 2048-bit modulus. Considered good protection for 128-bit keys. (ASA 9.0.1+ devices only).
- Diffie-Hellman Group 15: 3072-bit modulus. Considered good protection for 192-bit keys.
- Diffie-Hellman Group 16: 4096-bit modulus. Considered good protection for 256-bit keys.
- Diffie-Hellman Group 19: (256-bit elliptical curve field size). (ASA 9.0.1+ devices only).
- Diffie-Hellman Group 20: (384-bit elliptical curve field size). (ASA 9.0.1+ devices only).
- Diffie-Hellman Group 21: (521-bit elliptical curve field size). (ASA 9.0.1+ devices only).

- Diffie-Hellman Group 24: (2048-bit modulus and 256-bit prime order subgroup). (ASA 9.0.1+ devices only).

Related Topics

- Understanding IKE, page 26-5
- Configuring an IKE Proposal, page 26-9

Deciding Which Authentication Method to Use

Security Manager supports two methods for peer device authentication in a VPN communication:

- **Preshared Key**—Preshared keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. The same shared key must be configured at each peer or the IKE SA cannot be established.

  To use IKE successfully with this device authentication method, you must define various preshared key parameters. For more information, see the appropriate topic:

  - Site-to-site VPN, IKEv1 configuration—See Configuring IKEv1 Preshared Key Policies, page 26-48.
  - Site-to-site VPN, IKEv2 configuration—See Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.
  - Remote access IPsec VPN, IKEv1—Configured on the IPsec tab of the connection profile. See IPSec Tab (Connection Profiles), page 31-19.
  - Remote access IPsec VPN, IKEv2—You cannot use preshared keys when using IKEv2 in a remote access IPSec VPN. You must use certificates.

- **Certificate**—An authentication method in which RSA key pairs are used to sign and encrypt IKE key management messages. Certificates provide non-repudiation of communication between two peers, meaning that it can be proved that the communication actually took place. When using this authentication method, peers are configured to obtain digital certificates from a Certification Authority (CA). CAs manage certificate requests and issue certificates to participating IPsec network devices. These services provide centralized key management for the participating devices.

  While the use of preshared keys does not scale well, using a CA does improve the manageability and scalability of your IPsec network. With a CA, you do not need to configure keys between all encrypting devices. Instead, each participating device is registered with the CA, and requests a certificate from the CA. Each device that has its own certificate and the public key of the CA can authenticate every other device within a given CA's domain.

  To use IKE successfully with the Certificate authentication method, you must define parameters for CA authentication and enrollment. For more information, see the appropriate topic:

  - Site-to-site VPN, IKEv1 configuration—See Understanding Public Key Infrastructure Policies, page 26-51.
  - Site-to-site VPN, IKEv2 configuration—Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.
  - Remote access IPsec VPN, IKEv1—Configured on the IPsec tab of the connection profile as explained in IPSec Tab (Connection Profiles), page 31-19. You must also configure the Public Key Infrastructure policy with the same trustpoint; see Understanding Public Key Infrastructure Policies, page 26-51.

- Remote access IPsec VPN, IKEv2—Configure the global trustpoint on the IKEv2 Settings tab of the Global Settings policy as explained in Configuring VPN Global IKEv2 Settings, page 26-37. You must also configure the Public Key Infrastructure policy with the same trustpoint; see Understanding Public Key Infrastructure Policies, page 26-51.

Related Topics

- Understanding IKE, page 26-5
- Configuring an IKE Proposal, page 26-9

Configuring an IKE Proposal

In Security Manager, an IKE proposal is a mandatory policy when you configure a site-to-site or remote access IPsec VPN. When you use the configuration wizard to create a new IPsec VPN, an IKE Proposal policy is automatically assigned to the VPN; the policy might be the factory default, or it might be a shared policy specifically selected for the VPN. For more information about the IKE (Internet Key Exchange) key management protocol, see Understanding IKE, page 26-5.

Use the IKE Proposal policy to examine the current IKE proposals and to configure new proposals except for GET VPN topologies. For GET VPN, see Configuring the IKE Proposal for GET VPN, page 29-15.

Note Beginning with Cisco Security Manager version 4.17, you can configure and deploy IKE Proposal policy on ASA multi-context devices running the software version 9.9(2) or later.

Tips

- For site-to-site VPNs, you can select at most one IKE proposal per IKE version. For remote access IPsec VPNs, you can select multiple proposals for each IKE version; select all IKE proposals that are allowed in the remote access VPN.
- To configure IKEv2 (version 2), the device must be an ASA running ASA Software release 8.4(1) or higher.
- The IPsec Proposal policy must enable IKEv1, IKEv2, or both, to match the IKE proposals you configure in this policy. In cases where you cannot configure IKEv2 in the IPsec proposal, such as in Easy VPN topologies, IKEv2 is not supported. For more information, see Understanding IPsec Proposals, page 26-18.
- The IKEv1 Proposal objects specify whether preshared keys or certificates are used for authentication. If the IKEv1 Proposal object is of certificate authentication type, ensure that you specify the appropriate CA Server in the IKEv1 Public Key Infrastructure (from Policy Selector) policy. For preshared keys, ensure that the IKEv1 Preshared Key policy is assigned. For IKEv2, the object does not specify whether preshared keys or certificates are used, but other policies must define the authentication requirements. For more information, see Deciding Which Authentication Method to Use, page 26-8.
- For Regular IPSec VTI technology, you can specify only one of the IKE Proposals—IKEv1 proposal or IKEv2 proposal. That is, if you have selected IKE version 1 for the Regular IPSec VTI, (in the IKE Proposal window) you must be specifying the IKEv1 Proposal and leave the IKEv2 Proposal field blank and vice versa.

Related Topics

- Deciding Which Hash Algorithm to Use, page 26-6
- Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7

- Deciding Which Authentication Method to Use, page 26-8

Step 1 Do one of the following to open the IKE Proposal policy based on the type of VPN you are configuring:

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > IPSec VPN > IKE Proposal** from the Policy selector.
  - (Policy View) Select **Remote Access VPN > IPSec VPN > IKE Proposal** from the Policy Type selector. Select an existing policy or create a new one.
- For site-to-site VPNs, do one of the following:
  - Open the Site-to-Site VPN Manager Window, page 25-18, select a topology (other than GET VPN) in the VPNs selector, then select **IKE Proposal** in the Policies selector.
  - (Policy view) Select **Site-to-Site VPN > IKE Proposal** from the Policy Types selector. Select an existing shared policy or create a new one.

Step 2 Click **Select** against the appropriate IKE versions to choose the policy objects that define the settings for an IKE version 1 or version 2 proposal. Configure proposals only for those IKE versions supported in the VPN.

Note Beginning with 4.16, Cisco Security Manager does not support IKEv1 configuration for Firepower 9300 devices with distributed mode.

- To select an IKE proposal for site-to-site VPNs, simply highlight it in the available proposals list. For remote access IPsec VPNs, highlight the desired objects in the available proposals list and click **>>** to move them to the selected proposals list.
- To remove an IKE proposal for remote access IPsec VPNs, highlight it in the selected proposals list and click **<<** to move it to the available proposals list.
- To create a new IKE proposal, click the **Create (+)** button beneath the available proposals list. The Add IKEv1 or IKEv2 Proposal dialog box opens. For instructions on creating the object, see the following topics:
  - Configuring IKEv1 Proposal Policy Objects, page 26-10
  - Configuring IKEv2 Proposal Policy Objects, page 26-14
- To edit an object, or to view its settings, select it and click the **Edit (pencil)** button beneath the list.

Configuring IKEv1 Proposal Policy Objects

Use the IKEv1 Proposal dialog box to create, copy, and edit an IKEv1 proposal object.

Internet Key Exchange (IKE) version 1 proposal objects contain the parameters required for IKEv1 proposals when defining remote access and site-to-site VPN policies. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes security associations (SAs) for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. For more information about IKE proposals, see the following topics:

- Overview of IKE and IPsec Configurations, page 26-2
- Comparing IKE Version 1 and 2, page 26-4
- Understanding IKE, page 26-5
- Deciding Which Encryption Algorithm to Use, page 26-6
- Deciding Which Hash Algorithm to Use, page 26-6
- Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7
- Deciding Which Authentication Method to Use, page 26-8

Navigation Path

Select **Manage > Policy Objects**, then select **IKE Proposals > IKEv1 Proposals** from the Object Type Selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Tip You can also access this dialog box when configuring the IKE Proposal policy as explained in Configuring an IKE Proposal, page 26-9.

Related Topics

- Configuring IKEv2 Proposal Policy Objects, page 26-14
- Creating Policy Objects, page 6-9
- Policy Object Manager, page 6-4
- Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27

Field Reference

Table 26-1 IKEv1 Proposal Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object. A maximum of 1024 characters is allowed.
Priority	The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, the device tries to use the parameters defined in the policy with the next lowest priority number. Valid values range from 1 to 10000. The lower the number, the higher the priority. If you leave this field blank, Security Manager assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Table 26-1 IKEv1 Proposal Dialog Box (continued)

Element	Description
Encryption Algorithm	The encryption algorithm used to establish the Phase 1 SA for protecting Phase 2 negotiations: • AES-128—Encrypts according to the Advanced Encryption Standard using 128-bit keys. • AES-192—Encrypts according to the Advanced Encryption Standard using 192-bit keys. • AES-256—Encrypts according to the Advanced Encryption Standard using 256-bit keys. • DES—Encrypts according to the Data Encryption Standard using 56-bit keys. • 3DES—Encrypts three times using 56-bit keys. 3DES is more secure than DES, but requires more processing for encryption and decryption. It is less secure than AES. A 3DES license is required to use this option.
Hash Algorithm	The hash algorithm used in the IKE proposal. The hash algorithm creates a message digest, which is used to ensure message integrity. Options are: • SHA (Secure Hash Algorithm)—Produces a 160-bit digest. SHA is more resistant to brute-force attacks than MD5. • MD5 (Message Digest 5)—Produces a 128-bit digest. MD5 uses less processing time than SHA.

Table 26-1 IKEv1 Proposal Dialog Box (continued)

Element	Description
Modulus Group	The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Options are: **Tip** For IKEv1, ASA devices support groups 2 and 5 only. • 1—Diffie-Hellman Group 1 (768-bit modulus). **Note** DH group 1 is deprecated and will be removed in later ASA version. In later ASA versions, the default value will be Group 2. • 2—Diffie-Hellman Group 2 (1024-bit modulus). • 5—Diffie-Hellman Group 5 (1536-bit modulus, considered good protection for 128-bit keys, but group 14 is better). If you are using AES encryption, use this group (or higher). • 7—Diffie-Hellman Group 7 (163-bit elliptical curve field size). • 14—Diffie-Hellman Group 14 (2048-bit modulus, considered good protection for 128-bit keys). • 15—Diffie-Hellman Group 15 (3072-bit modulus, considered good protection for 192-bit keys). • 16—Diffie-Hellman Group 16 (4096-bit modulus, considered good protection for 256-bit keys). **Note** Although Diffie-Hellman groups 19-24 are listed, they are not supported for IKEv1 and will cause a validation error if selected for IKEv1 policies.
Lifetime	The lifetime of the security association (SA), in seconds. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. You can specify a value from 60 to 2147483647 seconds. The default is 86400.

Table 26-1 IKEv1 Proposal Dialog Box (continued)

Element	Description
Authentication Method	The method of authentication to use between the two peers. For information on how this selection determines which other policies you must configure, see Deciding Which Authentication Method to Use, page 26-8. Select one of the following: • Preshared Key—Preshared keys allow for a secret key to be shared between two peers and to be used by IKE during the authentication phase. If one of the participating peers is not configured with the same preshared key, the IKE SA cannot be established. • Certificate—An authentication method in which RSA key pairs are used to sign and encrypt IKE key management messages. This method provides non-repudiation of communication between two peers, meaning that it can be proved that the communication actually took place. When you use this authentication method, the peers are configured to obtain digital certificates from a Certification Authority (CA).
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Configuring IKEv2 Proposal Policy Objects

Use the IKEv2 Proposal dialog box to create, copy, and edit an IKEv2 proposal object. You can use IKEv2 proposals with ASA Software release 8.4(1)+ only.

Internet Key Exchange (IKE) version 2 proposal objects contain the parameters required for IKEv2 proposals when defining remote access and site-to-site VPN policies. IKE is a key management protocol that facilitates the management of IPsec-based communications. It is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and automatically establish IPsec security associations (SAs).

The IKE negotiation comprises two phases. Phase 1 negotiates a security association between two IKE peers, which enables the peers to communicate securely in Phase 2. During Phase 2 negotiation, IKE establishes security associations (SAs) for other applications, such as IPsec. Both phases use proposals when they negotiate a connection. Unlike IKEv1, in an IKEv2 proposal, you can select multiple algorithms and modulus groups from which peers can choose during the Phase 1 negotiation. For more information about IKE proposals, see the following topics:

- Overview of IKE and IPsec Configurations, page 26-2
- Comparing IKE Version 1 and 2, page 26-4
- Understanding IKE, page 26-5
- Deciding Which Encryption Algorithm to Use, page 26-6
- Deciding Which Hash Algorithm to Use, page 26-6
- Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7

Tip Unlike IKEv1, you do not specify the authentication method in the IKE proposal. For more information on how to configure the authentication method in IKEv2, see Deciding Which Authentication Method to Use, page 26-8.

Navigation Path

Select **Manage > Policy Objects**, then select **IKE Proposals > IKEv2 Proposals** from the Object Type Selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Tip You can also access this dialog box when configuring the IKE Proposal policy as explained in Configuring an IKE Proposal, page 26-9.

Related Topics

- Configuring IKEv1 Proposal Policy Objects, page 26-10
- Creating Policy Objects, page 6-9
- Policy Object Manager, page 6-4
- Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27

Field Reference

Table 26-2 IKEv2 Proposal Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object. A maximum of 1024 characters is allowed.
Priority	The priority value of the IKE proposal. The priority value determines the order of the IKE proposals compared by the two negotiating peers when attempting to find a common security association (SA). If the remote IPsec peer does not support the parameters selected in your first priority policy, the device tries to use the parameters defined in the policy with the next lowest priority number. Valid values range from 1 to 65535. The lower the number, the higher the priority. If you leave this field blank, Security Manager assigns the lowest unassigned value starting with 1, then 5, then continuing in increments of 5.

Table 26-2 IKEv2 Proposal Dialog Box (continued)

Element	Description
Encryption Algorithm	The encryption algorithm used to establish the Phase 1 SA for protecting Phase 2 negotiations. Click **Select** and select all of the algorithms that you want to allow in the VPN: • AES-GCM-256—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 256-bit keys. (ASA 5580 and ASA 5500-X Series devices running 9.0.1+ only). • AES-GCM-192—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 192-bit keys. (ASA 5580 and ASA 5500-X Series devices running 9.0.1+ only). • AES-GCM—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 128-bit keys. (ASA 5580 and ASA 5500-X Series devices running 9.0.1+ only). • AES-256—Encrypts according to the Advanced Encryption Standard using 256-bit keys. • AES-192—Encrypts according to the Advanced Encryption Standard using 192-bit keys. • AES—Encrypts according to the Advanced Encryption Standard using 128-bit keys. • 3DES—Encrypts three times using 56-bit keys. 3DES is more secure than DES, but requires more processing for encryption and decryption. It is less secure than AES. A 3DES license is required to use this option. • DES—Encrypts according to the Data Encryption Standard using 56-bit keys. • Null—No encryption algorithm.

Table 26-2 IKEv2 Proposal Dialog Box (continued)

Element	Description
Integrity (Hash) Algorithm	The integrity portion of the hash algorithm used in the IKE proposal. The hash algorithm creates a message digest, which is used to ensure message integrity. Click **Select** and select all of the algorithms that you want to allow in the VPN: **Note** If using AES-GCM, AES-GCM-192, or AES-GCM-256, you must select Null as the Integrity Algorithm. • SHA (Secure Hash Algorithm)—SHA is more resistant to brute-force attacks than MD5. Standard SHA produces a 160-bit digest. The following options, which are even more secure, are available for IKEv2 configurations on ASA 8.4(2+) devices: – SHA512—A 512-bit key. – SHA384—A 384-bit key. – SHA256—A 256-bit key. • MD5 (Message Digest 5)—Produces a 128-bit digest. MD5 uses less processing time than SHA. • Null—No encryption algorithm. For use with AES-GCM, AES-GCM-192, and AES-GCM-256 only.
Prf Algorithm	The pseudo-random function (PRF) portion of the hash algorithm used in the IKE proposal. In IKEv1, the Integrity and PRF algorithms are not separated, but in IKEv2, you can specify different algorithms for these elements. Click **Select** and select all of the algorithms that you want to allow in the VPN. The options are described above under Integrity Algorithm.

Table 26-2 IKEv2 Proposal Dialog Box (continued)

Element	Description
Modulus Group	The Diffie-Hellman group to use for deriving a shared secret between the two IPsec peers without transmitting it to each other. A larger modulus provides higher security but requires more processing time. The two peers must have a matching modulus group. Click **Select** and select all of the groups that you want to allow in the VPN: • 1—Diffie-Hellman Group 1 (768-bit modulus). **Note** DH group 1 is deprecated and will be removed in later ASA version. In later ASA versions, the default value will be Group 2. • 2—Diffie-Hellman Group 2 (1024-bit modulus). • 5—Diffie-Hellman Group 5 (1536-bit modulus, considered good protection for 128-bit keys, but group 14 is better). If you are using AES encryption, use this group (or higher). • 14—Diffie-Hellman Group 14 (2048-bit modulus, considered good protection for 128-bit keys). (ASA 9.0.1+ devices only). • 19—Diffie-Hellman Group 19 (256-bit elliptical curve field size). (ASA 9.0.1+ devices only). • 20—Diffie-Hellman Group 20 (384-bit elliptical curve field size). (ASA 9.0.1+ devices only). • 21—Diffie-Hellman Group 21 (521-bit elliptical curve field size). (ASA 9.0.1+ devices only). • 24—Diffie-Hellman Group 24 (2048-bit modulus and 256-bit prime order subgroup). (ASA 9.0.1+ devices only).
Lifetime	The lifetime of the security association (SA), in seconds. When the lifetime is exceeded, the SA expires and must be renegotiated between the two peers. As a general rule, the shorter the lifetime (up to a point), the more secure your IKE negotiations will be. However, with longer lifetimes, future IPsec security associations can be set up more quickly than with shorter lifetimes. You can specify a value from 120 to 2147483647 seconds. The default is 86400.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Understanding IPsec Proposals

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers, which can be devices in a site-to-site VPN or a device and user in remote access IPsec VPNs. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms called a transform set.

An IPsec proposal is used in Phase 2 of an IKE negotiation, as explained in Understanding IKE, page 26-5. The specific content of the proposal varies according to topology type (site-to-site or remote access) and device type, although the proposals are broadly similar and contain many of the same elements, such as IPsec transform sets.

The following topics explain IPsec proposal concepts and procedures in more detail:

- Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19
  - Understanding Crypto Maps, page 26-19
  - Understanding Transform Sets, page 26-20
  - Understanding Reverse Route Injection, page 26-21
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22
- Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40
- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3

Understanding IPsec Proposals for Site-to-Site VPNs

IPsec is one of the most secure methods for setting up a VPN. IPsec provides data encryption at the IP packet level, offering a robust security solution that is standards-based. Pure IPsec configurations cannot use routing protocols—the policy created is used for pure IPsec provisioning. You can configure pure IPsec on Cisco IOS routers, PIX Firewalls, Catalyst VPN Service Modules, and Adaptive Security Appliance (ASA) devices.

With IPsec, data is transmitted over a public network through tunnels. A tunnel is a secure, logical communication path between two peers. Traffic that enters an IPsec tunnel is secured by a combination of security protocols and algorithms called a transform set.

In Security Manager, you use an IPsec Proposal policy to define the settings required for a IPsec tunnels. An IPsec proposal is a collection of one or more crypto maps that are applied to the VPN interfaces on the devices. A crypto map combines all the components required to set up IPsec security associations, including transform sets. A crypto map can also be configured with Reverse Route Injection (RRI).

The following topics provide more information:

- Understanding Crypto Maps, page 26-19
- Understanding Transform Sets, page 26-20
- Understanding Reverse Route Injection, page 26-21

Related Topics

- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22

Understanding Crypto Maps

A crypto map combines all components required to set up IPsec security associations (SA), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA. A crypto map entry is a named series of CLI commands. Crypto map entries with the same crypto

map name (but different map sequence numbers) are grouped into a crypto map set, which is applied to the VPN interfaces on relevant devices. All IP traffic passing through the interface is evaluated against the applied crypto map set.

When two peers try to establish an SA, they must each have at least one compatible crypto map entry. The transform set defined in the crypto map entry is used in the IPsec security negotiation to protect the data flows specified by that crypto map's IPsec rules.

Dynamic crypto map policies are used in site-to-site VPNs when an unknown remote peer tries to initiate an IPsec security association with the local hub. The hub cannot be the initiator of the security association negotiation. Dynamic crypto policies allow remote peers to exchange IPsec traffic with a local hub even if the hub does not know the remote peer's identity. You can create a dynamic crypto policy on individual hubs or on a device group that contains hubs. The policy is written only to the hubs, not to any spokes that might be contained in the group. A dynamic crypto map policy essentially creates a crypto map entry without all the parameters configured. The missing parameters are later dynamically configured (as the result of an IPsec negotiation) to match a remote peer's requirements. The peer addresses for dynamic or static crypto maps are deduced from the VPN topology.

Dynamic crypto map policies apply only in a hub-and-spoke VPN configuration—in a point-to-point or full mesh VPN topology, you can apply only static crypto map policies.

Note (Site-to-site VPNs.) Except for Extranet VPNs, Security Manager can manage an existing VPN tunnel only if the tunnel's peers are managed by Security Manager. In such a case, Security Manager uses the same crypto map name for the tunnel on the peers. On subsequent deployments, only Security Manager tunnels are managed (Security Manager maintains a log of all tunnels that were configured).

Related Topics

- Understanding IPsec Proposals, page 26-18
- Understanding Transform Sets, page 26-20
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22

Understanding Transform Sets

A transform set is a combination of security protocols and algorithms that secure traffic in an IPsec tunnel. During the IPsec security association (SA) negotiation, peers search for a transform set that is the same at both peers. When such a transform set is found, it is applied to create an SA that protects data flows in the access list for that crypto map, protecting the traffic in the VPN.

There are separate IPsec transform sets for IKEv1 and IKEv2. With IKEv1 transform sets, for each parameter, you set one value. For IKEv2 transform sets, you can configure multiple encryption and integration algorithms for a single proposal. ASA devices order the settings from the most secure to the least secure and negotiate with the peer using that order. This allows you to potentially send a single proposal to convey all the allowed combinations instead of the need to send each allowed combination individually as with IKEv1.

You can specify a number of transform sets per IPsec proposal policy. If you are defining the policy on a spoke or a group of spokes, you do not usually have to specify more than one transform set. This is because the spoke's assigned hub would typically be a higher performance router capable of supporting any transform set that the spoke supports. However, if you are defining the policy on a hub for dynamic crypto, you should specify more than one transform set to ensure that there will be a transform set match between the hub and the unknown spoke. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security is used.

Security Manager provides predefined transform sets that you can use in your tunnel policies. You can also create your own transform sets. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27.

Selecting Tunnel Mode for IKEv1 Transform Sets

When defining an IKEv1 transform set, you must specify which IPsec mode of operation to use—tunnel mode or transport mode. You can use the AH and ESP protocols to protect an entire IP payload (Tunnel mode) or just the upper-layer protocols of an IP payload (Transport mode).

In tunnel mode (the default), the entire original IP datagram is encrypted, and it becomes the payload in a new IP packet. This mode allows a router to act as an IPsec proxy. That is, the router performs encryption on behalf of the hosts. The source's router encrypts packets and forwards them along the IPsec tunnel. The destination's router decrypts the original IP datagram and forwards it on to the destination system. The major advantage of tunnel mode is that the end systems do not need to be modified to enjoy the benefits of IPsec. Tunnel mode also protects against traffic analysis. With tunnel mode, an attacker can only determine the tunnel endpoints and not the true source and destination of the tunneled packets, even if they are the same as the tunnel endpoints.

In transport mode, only the IP payload is encrypted, and the original IP headers are left intact. This mode has the advantage of adding only a few bytes to each packet. It also allows devices on the public network to see the final source and destination of the packet. However, by passing the IP header in the clear, transport mode allows an attacker to perform some traffic analysis. For example, an attacker could see when a company's CEO sent many packets to another senior executive. However, the attacker would only know that IP packets were sent; the attacker would not be able to decipher the contents of the packets. With transport mode, the destination of the flow must be an IPsec termination device.

Note You cannot use transport mode for VPN topologies using regular IPsec or Easy VPN.

Related Topics

- Understanding IPsec Proposals, page 26-18
- Understanding Crypto Maps, page 26-19
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22

Understanding Reverse Route Injection

Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint. These protected hosts and networks are known as remote proxy identities. Each route is created on the basis of the remote proxy network and mask, with the next hop to this network being the remote tunnel endpoint. By using the remote VPN router as the next hop, the traffic is forced through the crypto process to be encrypted.

After the static route is created on the VPN router, this information is propagated to upstream devices, allowing them to determine the appropriate VPN router to which to send returning traffic in order to maintain IPsec state flows. This is particularly useful if multiple VPN routers are used at a site to provide load balancing or failover, or if the remote VPN devices are not accessible through a default route. Routes are created in either the global routing table or the appropriate virtual route forwarding (VRF) table.

Note Security Manager automatically configures RRI on devices with High Availability (HA) or on the IPsec Aggregator when VRF-Aware IPsec is configured. You can also configure RRI on a device's crypto map in a remote access VPN. See Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40 and Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3.

In Security Manager, the following options are available for configuring Reverse Route Injection:

- For dynamic crypto maps, routes are created upon the successful establishment of IPsec security associations (SAs) for those remote proxies. The next hop back to those remote proxies is through the remote VPN router whose address is learned and applied during the creation of the dynamic crypto map template. The routes are deleted after the SAs are deleted.
- The Remote Peer option (available for IOS devices only) enables you to specify an interface or address as the explicit next hop to the remote VPN device. Two routes are created. One route is the standard remote proxy ID and the next hop is the remote VPN client tunnel address. The second route is the actual route to the remote tunnel endpoint, when a recursive lookup is forced to impose that the remote endpoint is reachable via "next-hop." Creation of the second route for the actual next hop is very important for VRF-Aware IPsec when a default route must be overridden by a more explicit route.

  Note For devices using a VPN Services Module (VPNSM), the next hop is the interface or subinterface/VLAN on which the crypto map is applied.

- In the case of Remote Peer IP (available for IOS devices only), one route is created to a remote proxy by way of a user-defined next hop. The next hop can be used to override a default route to properly direct outgoing encrypted packets. This option reduces the number of routes created and supports those platforms that do not readily facilitate route recursion.

Related Topics

- Understanding IPsec Proposals, page 26-18
- Understanding Crypto Maps, page 26-19
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22

Configuring IPsec Proposals in Site-to-Site VPNs

Use the IPsec Proposal page to configure the IPsec proposal used during IKE Phase 2 negotiations for site-to-site VPN topologies with the exception of Easy VPN topologies.

IPsec proposals used with Easy VPN topologies, and with remote access VPNs, are significantly different than the basic site-to-site proposal explained in this topic. For information on IPsec proposals in these other topologies, see the following topics:

- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40
- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3

Navigation Path

- (Site-to-Site VPN Manager Window, page 25-18) Select a non-Easy VPN topology in the VPNs selector, then select **IPsec Proposal** in the Policies selector. If necessary, click the **IPsec Proposal** tab.
- (Policy view) Select **Site-to-Site VPN > IPsec Proposal** from the Policy Types selector. Select an existing shared policy or create a new one.

Related Topics

- Understanding IKE, page 26-5
- Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19

Field Reference

Table 26-3 IPsec Proposal Page, Site-to-Site VPNs (except Easy VPN)

Element	Description
Crypto Map Type (Hub and spoke and full mesh topologies only.)	A crypto map combines all the components required to set up IPsec security associations (SA). When two peers try to establish an SA, they must each have at least one compatible crypto map entry. For more information, see Understanding Crypto Maps, page 26-19. Select the type of crypto map you want to generate: • Static—Use a static crypto map in a point-to-point or full mesh VPN topology. • Dynamic—Dynamic crypto maps can only be used in a hub-and-spoke VPN topology. Dynamic crypto map policies allow remote peers to exchange IPsec traffic with a local hub, even if the hub does not know the remote peer's identity.
Enable IKEv1 Enable IKEv2	The IKE versions to use during IKE negotiations. IKEv2 is supported on ASA Software release 8.4(x) only. Similarly, beginning with 4.16, Cisco Security Manager does not support IKEv1 configurations for Firepower 9300 devices configured with distributed mode. Select either or both options as appropriate; you must select IKEv1 if any device in the topology does not support IKEv2. When you select both options in hub-and-spoke or full mesh topologies, Security Manager automatically assigns the IKE version to devices based on the OS type and version used by the device. You can change these assignments by clicking the **IKE Version** tab, then click the Select button beneath the IKEv1 Enabled Peers or IKEv2 Enabled Peers to change which version is assigned to the device. You can change the assignments for devices that support each version only; other devices are not selectable. For more information, see Selecting the IKE Version for Devices in Site-to-Site VPNs, page 26-26.

Table 26-3 IPsec Proposal Page, Site-to-Site VPNs (except Easy VPN) (continued)

Element	Description
Transform Sets IKEv2 Transform Sets	The transform sets to use for your tunnel policy. Transform sets specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel. The transform sets are different for each IKE version; select objects for each supported version. You can select up to 11 transform sets for each. For more information, see Understanding Transform Sets, page 26-20. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security will be used. Click **Select** to select the IPsec transform set policy objects to use in the topology. If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27. **Note** IKEv1 Transform sets can use tunnel mode or transport mode of IPsec operation. However, you cannot use transport mode in IPsec or Easy VPN topologies.
Enable Perfect Forward Secrecy Modulus Group	Whether to use Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. The unique session key protects the exchange from subsequent decryption, even if the entire exchange was recorded and the attacker has obtained the preshared or private keys used by the endpoint devices. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the **Modulus Group** list. For an explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7. **Note** DH group 1 is deprecated and will be removed in later ASA version. In later ASA versions, the default value will be Group 2.
Lifetime (sec) Lifetime (kbytes)	The global lifetime settings for the crypto IPsec security association (SA). You can specify the IPsec lifetime in seconds, in kilobytes, or both. • Seconds (sec)—The number of seconds an SA will exist before expiring. The default is 3600 seconds (one hour). • Kilobytes (kbytes)—The volume of traffic (in kilobytes) that can pass between IPsec peers using a given SA before it expires. Valid values depend on the device type. Enter a value within the range 10-2147483647 for an IOS router, and 2560-536870912 for an ASA/PIX7.0+ device. The default value is 4,608,000 kilobytes.

Table 26-3 IPsec Proposal Page, Site-to-Site VPNs (except Easy VPN) (continued)

Element	Description
QoS Preclassify	Supported on Cisco IOS routers, except 7600 devices. When selected, enables the classification of packets before tunneling and encryption occur. The Quality of Service (QoS) for VPNs feature enables Cisco IOS QoS services to operate with tunneling and encryption on an interface. The QoS features on the output interface classify packets and apply the appropriate QoS service before the data is encrypted and tunneled, enabling traffic flows to be adjusted in congested environments, and resulting in more effective packet tunneling.
Reverse Route	Supported on ASA devices, PIX 7.0+ devices, and Cisco IOS routers except 7600 devices. Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint. For more information, see Understanding Reverse Route Injection, page 26-21. Select one of the following options to configure RRI on the crypto map: • **None**—Disables the configuration of RRI on the crypto map. • **Standard**—(ASA, PIX 7.0+, IOS devices) Creates routes based on the destination information defined in the crypto map access control list (ACL). This is the default option. • **Remote Peer**—(IOS devices only) Creates two routes, one for the remote endpoint and one for route recursion to the remote endpoint via the interface to which the crypto map is applied. • **Remote Peer IP**—(IOS devices only) Specifies an address as the explicit next hop to the remote VPN device. Enter the IP address or a network/host object that specifies the address, or click **Select** to select the network/host object from a list or to create a new object. **Note** If you use network/host objects, you can select the **Allow Value Override per Device** option in the object to override the IP address, if required, for specific devices that use this object.
Enable Dynamic RRI	**Note** This option is supported from ASA 9.7(1) onwards. It is only applicable if IKEV2 is enabled or Static Crypto Maps are selected. When enabled, the crypto map does not install the reverse-route during configuration but defers it till the IPsec security associations (SA) come up.
ESPv3 Settings (ASA 9.0.1+ only) Specify whether incoming ICMP error messages are validated for cryptography and dynamic cryptography maps, set the per-security association policy, or enable traffic flow packets:	
Validate Incoming ICMP error messages	Whether to validate those ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.

Table 26-3 IPsec Proposal Page, Site-to-Site VPNs (except Easy VPN) (continued)

Element	Description
Enable Do Not Fragment (DF) Policy	Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header. Choose one of the following: • Set—Sets and uses the DF bit. • Copy—Maintains the DF bit. • Clear—Ignores the DF bit.
Enable Traffic Flow Confidentiality (TFC) Packets	Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. **Note** You must have an IKEv2 IPsec proposal set on the Tunnel Policy (Crypto Map) Basic tab before enabling TFC. Traffic Flow Confidentiality is not available when IKEv1 is enabled. Use the Burst, Payload Size, and Timeout parameters to generate random length packets at random intervals across the specified SA.

Selecting the IKE Version for Devices in Site-to-Site VPNs

Use the IKE Version tab in the IPsec Proposal page to select which version of IKE to use for each device in a hub-and-spoke or full mesh site-to-site VPN. This tab appears only in the Site-to-Site VPN Manager; you cannot configure the options in Policy view, because they are specific to the actual devices in a VPN topology.

The IKE Version tab contains two lists: IKEv1 Enabled Peers and IKEv2 Enabled Peers. When you configure the IPsec proposal, as described in Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22, you select which IKE versions to allow in the VPN (version 1, version 2, or both). Security Manager automatically chooses which IKE version to use for a device based on the OS version used by the device. For example, IOS routers always appear in the IKEv1 Enabled Peers list. If a device supports both IKEv1 and IKEv2, it appears in both lists.

You need to alter the selection only if you are allowing both IKE versions in a VPN and you want to specifically prevent some IKEv2-capable devices from using one of the IKE versions.

To change which IKE version is allowed for a device, click the **Select** button beneath the list from which you want to remove the device (or to add the device after previously removing it). A selection dialog box opens where you can do the following (click **OK** to confirm your choices):

- To remove a device, so that it cannot use the IKE version, highlight it in the Selected Peers list and click **<<** to move it to the Available Peers list.
- To add a device, so that it is allowed to use the IKE version, highlight it in the Available Peers list and click **>>** to move it to the Selected Peers list.

Tip The selection lists include only those devices that support both IKE versions, because you cannot change the version selection for devices that support a single version. IKEv2 is supported on ASA Software 8.4(1)+.

Navigation Path

(Site-to-Site VPN Manager Window, page 25-18) Select a non-Easy VPN topology in the VPNs selector, then select **IPsec Proposal** in the Policies selector. Click the **IKE Version** tab.

Related Topics

- Understanding IKE, page 26-5
- Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19

Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects

Use the Add or Edit IPSec Transform Set dialog box to configure IPSec transform sets for use in IKE negotiations.

You can create IPSec transform set objects for use in IPSec proposals when defining IPSec-protected traffic in site-to-site and remote access VPNs. During IPSec security association negotiation, the peers agree to use a particular transform set when protecting a particular data flow.

Two different security protocols are included within the IPSec standard:

- Encapsulating Security Protocol (ESP)—Provides authentication, encryption, and anti-replay services. ESP is IP protocol type 50.
- Authentication Header (AH)—Provides authentication and anti-replay services. AH does not provide encryption and has largely been superseded by ESP. It is also supported on routers only. AH is IP protocol type 51.

Note We recommend using both encryption and authentication on IPSec tunnels.

There are separate IPsec transform set objects based on the IKE version, IKEv1 or IKEv2:

- When you create an IPSec IKEv1 transform set object, you select the mode in which IPSec should operate, as well as define the required encryption and authentication types. Additionally, you can select whether to include compression in the transform set. You can select single options for the algorithms, so if you want to support multiple combinations in a VPN, you must create multiple IPsec IKEv1 transform set objects.
- When you create an IPSec IKEv2 transform set object, you can select all of the encryption and hash algorithms that you will allow in a VPN. During IKEv2 negotiations, the peers select the most appropriate options that each support.

Note If you configure an IPsec IKEv1 or IKEv2 Proposal on a device, you must use the configured Proposal for that device. For example, in a Site-to-Site (Point-to-Point) VPN configuration, the endpoint (interface) configured with the IPsec Proposal can be used in generating the crypto map. However, if the configured Proposal is not used by Security Manager for that device, in the following preview configuration, Security Manager will generate a negate command and the configured IPsec Proposal will be negated by Security Manager.

Navigation Path

Select **Manage > Policy Objects**, then select **IPSec Transform Sets > IPSec IKEv1 Transform Sets** or **IPSec Transform Sets > IPSec IKEv2 Transform Sets** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Transform Sets, page 26-20
- Overview of IKE and IPsec Configurations, page 26-2
- Comparing IKE Version 1 and 2, page 26-4
- Understanding IKE, page 26-5
- Understanding IPsec Proposals, page 26-18
- IPsec Proposal Editor (ASA, PIX 7.0+ Devices), page 31-41
- IPsec Proposal Editor (IOS, PIX 6.3 Devices), page 33-4
- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40
- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3
- Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring IKEv1 Proposal Policy Objects, page 26-10
- Creating Policy Objects, page 6-9
- Policy Object Manager, page 6-4

Field Reference

Table 26-4 IPSec IKEv1 or IKEv2 Transform Set Dialog Box

Element	Description
Name	The name of the policy object. A maximum of 128 characters is allowed.
Description	A description of the policy object. A maximum of 1024 characters is allowed.
Mode (IKEv1 only.)	The mode in which the IPSec tunnel operates: • Tunnel—Tunnel mode encapsulates the entire IP packet. The IPSec header is added between the original IP header and a new IP header. This is the default. Use tunnel mode when the firewall is protecting traffic to and from hosts positioned behind the firewall. Tunnel mode is the normal way regular IPSec is implemented between two firewalls (or other security gateways) that are connected over an untrusted network, such as the Internet. • Transport—Transport mode encapsulates only the upper-layer protocols of an IP packet. The IPSec header is inserted between the IP header and the upper-layer protocol header (such as TCP). Transport mode requires that both the source and destination hosts support IPSec, and can only be used when the destination peer of the tunnel is the final destination of the IP packet. Transport mode is generally used only when protecting a Layer 2 or Layer 3 tunneling protocol such as GRE, L2TP, and DLSW.

Table 26-4 IPSec IKEv1 or IKEv2 Transform Set Dialog Box (continued)

Element	Description
ESP Encryption	The Encapsulating Security Protocol (ESP) encryption algorithm that the transform set should use. For more information on the following options, see Deciding Which Encryption Algorithm to Use, page 26-6. For IKEv1, select one of the following options. For IKEv2, click **Select** to open a dialog box where you can select all of the options you want to support: **Note** AES-GCM/GMAC can only be configured on 5580 and newer ASA platforms. • (Blank)—Do not use ESP encryption. • DES—Encrypts according to the Data Encryption Standard using 56-bit keys. • 3DES—Encrypts three times using 56-bit keys. 3DES is more secure than DES, but requires more processing for encryption and decryption. It is less secure than AES. A 3DES license is required to use this option. • AES-128 (AES)—Encrypts according to the Advanced Encryption Standard using 128-bit keys. • AES-192—Encrypts according to the Advanced Encryption Standard using 192-bit keys. • AES-256—Encrypts according to the Advanced Encryption Standard using 256-bit keys. • ESP-Null (NULL)—A null encryption algorithm. Transform sets defined with ESP-Null provide authentication without encryption; this is typically used for testing purposes only. • AES-GCM (IKEv2 only)—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 128-bit keys. (ASA 9.0.1+ devices only). • AES-GCM-192 (IKEv2 only)—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 192-bit keys. (ASA 9.0.1+ devices only). • AES-GCM-256 (IKEv2 only)—Encrypts according to the Advanced Encryption Standard in Galois/Counter Mode using 256-bit keys. (ASA 9.0.1+ devices only). • AES-GMAC (IKEv2 only)—Encrypts according to the Advanced Encryption Standard Galois Message Authentication Codeusing 128-bit keys. • AES-GMAC-192 (IKEv2 only)—Encrypts according to the Advanced Encryption Standard Galois Message Authentication Codeusing 192-bit keys. • AES-GMAC-256 (IKEv2 only)—Encrypts according to the Advanced Encryption Standard Galois Message Authentication Code using 256-bit keys.

Table 26-4 IPSec IKEv1 or IKEv2 Transform Set Dialog Box (continued)

Element	Description
ESP Hash Algorithm (IKEv1) ESP Integration Algorithm (IKEv2) AH Hash Algorithm (IKEv1 only)	The hash or integrity algorithm to use in the transform set for authentication. For IKEv1, the default is to use SHA for ESP authentication and to not use AH authentication. For IKEv2, there is no default. The AH hash algorithm is used on routers only. For IKEv1, select one of the following options. For IKEv2, click **Select** to open a dialog box where you can select all of the options you want to support. • None—Does not perform ESP or AH authentication. • SHA, SHA-1 (Secure Hash Algorithm version 1)—Produces a 160-bit digest. SHA is more resistant to brute-force attacks than MD5, but requires more processing time. The following options, which are even more secure, are available for IKEv2 configurations on ASA 8.4(2+) devices: – SHA512—A 512-bit key. – SHA384—A 384-bit key. – SHA256—A 256-bit key. • MD5 (Message Digest 5)—Produces a 128-bit digest. MD5 uses less processing time than SHA, but is less secure. • Null—No encryption algorithm. For use with AES-GCM, AES-GCM-192, AES-GCM-256, AES-GMAC, AES-GMAC-192, and AES-GMAC-256 only.
Compression (IKEv1 only, IOS devices only.)	Whether to compress the data in the IPSec tunnel using the Lempel-Ziv-Stac (LZS) algorithm.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Configuring VPN Global Settings

You can define global settings that apply to all devices in your remote access or site-to-site VPN topology. These settings include Internet Key Exchange (IKE), IKEv2, IPsec, NAT, and fragmentation definitions. The global settings typically have defaults that work in most situations, so configuring the Global Settings policy is optional in most cases; configure it only if you need non-default behavior or if you are supporting IKEv2 negotiations in a remote access IPsec VPN.

Note The VPN Global Settings policy for site-to-site VPNs applies to all technologies except GET VPN. For an explanation of global settings for GET VPN, see Configuring Global Settings for GET VPN, page 29-16.

Step 1 Do one of the following to open the global settings policy based on the type of VPN you are configuring:

- For remote access VPNs, do one of the following:

  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one.
- For site-to-site VPNs, do one of the following:
  - Open the Site-to-Site VPN Manager Window, page 25-18, select a topology in the VPNs selector, then select **VPN Global Settings** in the Policies selector.
  - (Policy view) Select **Site-to-Site VPN > VPN Global Settings** from the Policy Types selector. Select an existing shared policy or create a new one.

Step 2 Select the desired tab and configure the settings as needed:

- **ISAKMP/IPsec Settings**—To configure global settings for IKE and IPsec. For detailed information about the options, see Configuring VPN Global ISAKMP/IPsec Settings, page 26-33.
- **IKEv2 Settings**—To configure global settings for IKE version 2 negotiations. For detailed information about the options, see Configuring VPN Global IKEv2 Settings, page 26-37.
- **NAT Settings**—To configure NAT behavior. For detailed information about the options, see Configuring VPN Global NAT Settings, page 26-42. Also see Understanding NAT in VPNs, page 26-41.
- **Address Assignment**—To specify one or more methods of address assignment to remote clients, see Configuring VPN Global Address Assignment Settings, page 26-31. The address assignment applies only to remote access VPN.
- **General Settings**—To configure fragmentation behavior and some other miscellaneous options. For detailed information about the options, see Configuring VPN Global General Settings, page 26-44.

Configuring VPN Global Address Assignment Settings

Use the Address Assignment tab of the VPN Global Settings page to specify one or more methods of address assignment to remote clients. The available methods are:

- Obtain IP addresses from an authentication server.
- Obtain IP addresses from a DHCP server.
- Obtain IP addresses from an internally configured pool.

Note You can configure address assignment on devices running the ASA software version 7.0(1) or higher. By default all the methods are enabled.

Address Assignment is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector. Click the **Address Assignment** tab.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one, then click the **Address Assignment** tab.

Related Topics

- Configuring VPN Global Settings, page 26-30

Field Reference

Table 26-5 VPN Global Settings Page, Address Assignment Tab

Element	Description
IPv4 Address Assignment Priority	
Use Authentication Server	Choose to assign IPv4 addresses retrieved from an authentication server on a per-user basis. If you are using an authentication server (external or internal) that has IPv4 addresses configured, we recommend using this method. If you select this option, use the **Platform > Device Admin > AAA** policy to define the AAA server groups to use for authenticating user access. This method is available for IPv4 and IPv6 assignment policies.
Use DHCP	Choose to obtain IP addresses from a DHCP server. If you use DHCP, you must configure the server using the **Platform > Device Admin > Server Access > DHCP Server** from the Device Policy selector. You must also define the range of IP addresses that the DHCP server can use. This method is available for IPv4 assignment policies.
Use internal address pools	Choose to have the ASA assign IPv4 addresses from an internally configured pool. Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, you must configure the IP address pools. To configure IP address pools, on the Device view, select **NAT > Address Pools** from the Device Policy Selector. Or on the Policy view, select **NAT (PIX/ASA/FWSM) > Address Pools** from the Policy Type selector, then select an existing policy from the Shared Policy selector, or right-click **Address Pools** to create a new policy.
Allow the reuse of an IP address - minutes after it is released	Delays the reuse of an IP address after its return to the address pool. Adding a delay helps to prevent problems firewalls can experience when an IP address is reassigned quickly. By default, this is unchecked, meaning the ASA does not impose a delay. To add a delay, check the box and enter the number of minutes in the range 0 - 480 to delay IP address reassignment. **Note** This feature is available on devices running the ASA software version 8.0(3) or higher.
IPv6 Address Assignment Priority - Security Manager 4.12 onwards for ASA devices running version 9.0 or later.	
Use Authentication Server	Choose to assign IPv6 addresses retrieved from an authentication server on a per-user basis. If you are using an authentication server (external or internal) that has IPv6 addresses configured, we recommend using this method. If you select this option, use the **Platform > Device Admin > AAA** policy to define the AAA server groups to use for authenticating user access.

Table 26-5 VPN Global Settings Page, Address Assignment Tab (continued)

Element	Description
Use internal address pools	Choose to have the ASA assign IPv6 addresses from an internally configured pool. Internally configured address pools are the easiest method of address pool assignment to configure. If you use this method, you must configure the IP address pools. To configure IPv6 address pools, on the Device view, select **NAT > Address Pools** from the Device Policy Selector. Or on the Policy view, select **NAT (PIX/ASA/FWSM) > Address Pools** from the Policy Type selector, then select an existing policy from the Shared Policy selector, or right-click **Address Pools** to create a new policy.

Configuring VPN Global ISAKMP/IPsec Settings

Use the ISAKMP/IPsec Settings tab of the VPN Global Settings page to specify global settings for Internet Key Exchange (IKE) and IPsec.

The Internet Key Exchange (IKE) protocol, also called the Internet Security Association and Key Management Protocol (ISAKMP) is the negotiation protocol that lets two hosts agree on how to build an IPsec security association. Each ISAKMP negotiation is divided into a Phase 1 and Phase 2. Phase 1 creates the first tunnel, which protects ISAKMP negotiation messages. Phase 2 creates the tunnel that protects data.

To set terms for ISAKMP negotiations, you create an IKE proposal. For more information, see Configuring an IKE Proposal, page 26-9.

About IKE Keepalive

With IKE keepalive, tunnel peers exchange messages that demonstrate they are available to send and receive data in the tunnel. Keepalive messages transmit at set intervals, and any disruption in that interval results in the creation of a new tunnel, using a backup device.

Devices that rely on IKE keepalive for resiliency transmit their keepalive messages regardless of whether they are exchanging other information. These keepalive messages can therefore create a small but additional demand on your network.

A variation on IKE keepalive called keepalive dead-peer detection (DPD) sends keepalive messages between peer devices only when no incoming traffic is received and outbound traffic needs to be sent. If you want to send DPD keepalive messages when no incoming traffic is received regardless of whether there is any outbound traffic, you can specify this using the Periodic option.

Navigation Path

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector. Click the **ISAKMP/IPsec Settings** tab.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one, then click the **ISAKMP/IPsec Settings** tab.
- For site-to-site VPNs, do one of the following:

- Open the Site-to-Site VPN Manager Window, page 25-18, select a topology in the VPNs selector, then select **VPN Global Settings** in the Policies selector. Click the **ISAKMP/IPsec Settings** tab.
- (Policy view) Select **Site-to-Site VPN > VPN Global Settings** from the Policy Types selector. Select an existing shared policy or create a new one, then click the **ISAKMP/IPsec Settings** tab.

Related Topics

- Configuring VPN Global Settings, page 26-30
- Understanding IKE, page 26-5
- Understanding IPsec Proposals, page 26-18

Field Reference

Table 26-6 VPN Global Settings Page, ISAKMP/IPsec Settings Tab

Element	Description
ISAKMP Settings	
Enable Keepalive	Whether to configure dead-peer detection (DPD) settings. If the peer fails to respond, a new tunnel is constructed on the assumption that the peer is no longer available. IKE keepalive is defined on the spokes in a hub-and-spoke VPN topology, on both devices in a point-to-point VPN topology, or in remote access VPN configurations. Configure the following options: • Interval—The number of seconds the peer can be idle before beginning keepalive monitoring. The range is 10-3600 seconds. The default is 10, although the ASA device default for remote access groups is 300. • Retry—The interval in seconds between retries after a keepalive response has not been received. The range is 2-10 seconds for ASA, 2-60 for IOS devices. The default is 2 seconds. • Periodic—(Routers running IOS Software version 12.3(7)T and later, except 7600 devices.) Whether to send DPD keepalive messages at regular intervals regardless of IPsec traffic. This changes how the interval value is used. • Infinite—(ASA only.) Whether to ignore the interval and retry settings and allow the peer to be idle indefinitely.

Table 26-6 VPN Global Settings Page, ISAKMP/IPsec Settings Tab (continued)

Element	Description
Identity	During Phase I IKE negotiations, peers must identify themselves to each other. Select one of the following: • Address—Use the IP address of the host exchanging ISAKMP identity information. This is the default. • Hostname—Use the fully-qualified domain name of the host exchanging ISAKMP identity information. • Auto/DN—Use automatic selection or distinguished name based on device type: – Distinguished Name (IOS devices only)—Use a distinguished name (DN) to identify a user group name. – Auto (ASA devices only)—Determine ISAKMP negotiation by connection type; IP address for preshared key or certificate distinguished name for certificate authentication.
SA Requests System Limit	Supported on routers running Cisco IOS Software Release 12.3(8)T and later, except 7600 routers. The maximum number of SA requests allowed before IKE starts rejecting them, from 0 to 99999. The number must equal or exceed the number of peers, or the VPN tunnels might be disconnected.
SA Requests System Threshold	Supported on Cisco IOS routers and Catalyst 6500/7600 devices. The percentage of system resources that can be used before IKE starts rejecting new SA requests. The default is 75 percent.
Enable Aggressive Mode (Site to site VPNs only.)	Supported on ASA devices and PIX 7.0+ devices. When selected, enables you to use aggressive mode in ISAKMP negotiations. Aggressive mode is enabled by default.
IPsec Settings	
Enable Lifetime	Select to enable you to configure the global lifetime settings for the crypto IPsec security associations (SAs) on the devices in your site-to-site or remote access VPN. Configure the following: • Lifetime (secs)—The number of seconds a security association will exist before expiring. The default is 3,600 seconds (1 hour). • Lifetime (kbytes)—The volume of traffic (in kilobytes) that can pass between IPsec peers using a given security association before it expires. The default is 4,608,000 kilobytes.

Table 26-6 VPN Global Settings Page, ISAKMP/IPsec Settings Tab (continued)

Element	Description
Xauth Timeout	Supported on Cisco IOS routers and Catalyst 6500/7600 devices in remote access VPN and Easy VPN topologies only. The number of seconds the device will wait for a system response to the Xauth challenge. When negotiating tunnel parameters for establishing IPsec tunnels in a remote access or Easy VPN configuration, Xauth adds another level of authentication that identifies the user who requests the IPsec connection. Using the Xauth feature, the client waits for a username/password (Xauth) challenge after the IKE SA has been established. When the end user responds to the challenge, the response is forwarded to the IPsec peers for an additional level of authentication.
Max Sessions	Supported on ASA devices and PIX 7.0+ devices. The maximum number of security associations (SAs) that can be enabled simultaneously on the device. The maximum number differs based on device model. For ASA devices, the limits are: • 5505—10 sessions. • 5510—250 sessions. • 5520—750 sessions. • 5540, 5550, 5585-X with SSP-10—5000 sessions. • 5580, 5585-X (other models)—10000 sessions.
Enable IPsec via Sysopt	Supported on ASA devices, and PIX Firewalls versions 6.3 or 7.0+. Whether to bypass the access rules defined on the VPN interface for VPN traffic. By default, the device allows VPN traffic to terminate on an interface; you do not need to allow IKE or ESP (or other types of VPN packets) in an interface access list. By default, you also do not need an interface access list for local IP addresses of decrypted VPN packets. Because the VPN tunnel was terminated successfully using VPN security mechanisms, this feature simplifies configuration and maximizes the device performance without any security risks. (Group policy and per-user authorization access lists still apply to the traffic.) If you deselect this option, the interface access rules are also applied to VPN traffic. The access list applies to the local IP address and not to the original client IP address used before the VPN packet was decrypted. The command applied is **no sysopt connection permit-vpn**.
Enable IPsec inner routing lookup (Security Manager version 4.12 onwards for ASA devices 9.6(2) or later)	To enable per-packet routing lookups for the IPsec inner packets. This checkbox is deselected by default.

Table 26-6 VPN Global Settings Page, ISAKMP/IPsec Settings Tab (continued)

Element	Description
Enable SPI Recovery (Site-to-site VPNs only.)	Supported on routers running IOS version 12.3(2)T and later, in addition to Catalyst 6500/7600 devices running version 12.2(18)SXE and later. When selected, enables the SPI recovery feature to configure your device so that if an invalid SPI (Security Parameter Index) occurs, an IKE SA will be initiated. SPI is a number which, together with a destination IP address and security protocol, uniquely identifies a particular security association. When using IKE to establish security associations, the SPI for each security association is a pseudo-randomly derived number. Without IKE, the SPI is manually specified for each security association. When an invalid SPI occurs during IPsec packet processing, the SPI recovery feature enables an IKE SA to be established.
ESPv3 Settings	
Enable PMTU (Path Maximum Transmission Unit) Aging	Supported for IKEv2 on ASA devices versions 9.0.1+. Whether to enable Path Maximum Transmission Unit aging. If you select this option, configure the interval, in minutes, at which the PMTU value is reset to its original value. The value can be from 10 to 30 minutes. The default is 10 minutes.

Configuring VPN Global IKEv2 Settings

Use the IKEv2 Settings tab of the VPN Global Settings page to specify global settings for Internet Key Exchange (IKE) version 2. These settings apply to ASA 8.4(x) devices only.

Internet Key Exchange (IKE), also called Internet Security Association and Key Management Protocol (ISAKMP), is the negotiation protocol that lets two hosts agree on how to build an IPsec security association (SA).

Preventing DoS Attacks by Limiting IKEv2 Open SAs

You can prevent denial-of-service (DoS) attacks for IPsec IKEv2 connections by always cookie challenging incoming security associations (SAs) or by limiting the number of open SAs and cookie challenging any additional connections. By default, the ASA does not limit the number of open SAs and never cookie challenges SAs.

You can also limit the number of SAs allowed, which stops further connections from negotiating to protect against memory or CPU attacks that the cookie-challenge feature may be unable to thwart. Limiting the maximum number of SAs can protect the current connections.

With a DoS attack, an attacker initiates the attack when the peer device sends an SA initiate packet and the ASA sends its response, but the peer device does not respond further. If the peer device does this continually, all the allowed SA requests on the ASA can be used up until it stops responding.

Enabling a threshold percentage for cookie challenges limits the number of open SA negotiations. For example, with the default setting of 50%, when 50% of the allowed SAs are in-negotiation (open), the ASA cookie challenges any additional SA initiate packets that arrive. For the Cisco ASA 5580 with 10,000 allowed IKEv2 SAs, after 5000 SAs become open, any more incoming SAs are cookie-challenged.

If used in conjunction with the **Maximum SAs in Negotiation** option, configure a lower cookie-challenge threshold.

Navigation Path

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector. Click the **IKEv2 Settings** tab.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one, then click the **IKEv2 Settings** tab.
- For site-to-site VPNs, do one of the following:
  - Open the Site-to-Site VPN Manager Window, page 25-18, select a topology in the VPNs selector, then select **VPN Global Settings** in the Policies selector. Click the **IKEv2 Settings** tab.
  - (Policy view) Select **Site-to-Site VPN > VPN Global Settings** from the Policy Types selector. Select an existing shared policy or create a new one, then click the **IKEv2 Settings** tab.

Related Topics

- Configuring VPN Global Settings, page 26-30
- Understanding IKE, page 26-5
- Understanding IPsec Proposals, page 26-18
- Configuring Cluster Load Balance Policies (ASA), page 31-5

Field Reference

Table 26-7 VPN Global Settings Page, IKEv2 Settings Tab

Element	Description
Maximum SAs	The number of allowed IKEv2 connections (security associations) on the device. The default limit is the maximum number of connections specified by the device license, which differs by device model. Specify a number only if you want to create a limit that is lower than the device license. The range is 1 to 10000.
Maximum SAs in Negotiation	The maximum number of IKEv2 security associations (SAs) that can be in negotiation at any time as a percentage of the maximum allowed SAs. The default is no limit on SAs in negotiation, so it is possible for all available SAs to be in negotiation. The range is 1 to 100%. If you configure this option and also enable custom cookie challenge, configure the cookie challenge threshold lower than this limit.

Table 26-7 VPN Global Settings Page, IKEv2 Settings Tab (continued)

Element	Description
Enable Cookie Challenge	Whether to send cookie challenges to peer devices in response to SA initiate packets, which can help thwart denial of service (DoS) attacks. The default is to use cookie challenges when 50% of the available SAs are in negotiation. Select one of these options: • **Custom**—Cookie challenge when the number of SAs in negotiation exceeds the total number of allowed SAs on the device based on percentage (SAs in negotiation as a percentage of total allowed SAs). In Custom Cookie Challenge, enter the percentage that triggers cookie challenges for any future SA negotiations. The range is 1 to 100%. The default is 50%. • **Never**—The device never uses cookie challenge. • **Always**—The device always uses cookie challenge, regardless of the percentage of SAs in negotiation.
Remote Access Authentication RA Trustpoint (Remote access VPN only.)	(Required when supporting IKEv2 negotiations.) The PKI enrollment object that identifies the Certificate Authority (CA) server that the device can use to authenticate itself to the remote user. This authorization is required before the user can select a connection profile and log into the VPN. This CA server is used in remote access IKEv2 IPsec VPNs only. Click **Select** to select the object or to create a new one. **Note** Beginning with Cisco Security Manager version 4.17, you can configure remote access authentication on ASA 9.9(2) or later multi-context devices. **Tip** You must also select this PKI enrollment object in the Remote Access VPN > Public Key Infrastructure policy.

Table 26-7 VPN Global Settings Page, IKEv2 Settings Tab (continued)

Element	Description
Load Balancing Settings Redirect Connections During (Remote access VPNs only.)	If you configure load balancing, using the ASA Cluster Load Balance policy, you can specify the IKEv2 negotiation phase in which a user can be redirected to another device in the cluster. Select one of these options: • **INIT**—Redirect at unauthenticated initiation requests (the first IKEv2 message IKE_SA_INIT), before any group or user authentication. – Pros—This option allows the primary server to do minimal processing and state keeping (using CPU and memory) prior to redirecting the connection. – Cons—This option is not as secure as AUTH (even though security risks are minimal) because anyone can get a redirected IP address without authenticating at all. • **AUTH** (the default)—Redirect during authentication (during IKE_AUTH). The device still has not identified or authenticated the user at this point, but it allows the client to authenticate the server to make sure it can trust the redirection that it receives. – Pros—This option is more secure as the reply is encrypted in the IKEv2 tunnel and it allows the client side to authenticate the server before retrying with the redirected IP address, providing better DoS protection than the INIT option. – Cons—This option requires more processing as the IKEv2 tunnel needs to be almost brought up before redirecting, although child SAs and data tunnels do not need to be brought up. The client is not authenticated at all. Note that IKEv1 redirection occurs after group authentication of both sides of the tunnel.
Enable Invalid Selectors Notification	Whether to enable sending an IKE notification to the peer when an inbound packet is received on an SA that does not match the traffic selectors for that SA. This feature is available in Security Manager version 4.9 onwards for ASA devices version 9.4(1) or higher.
Fragmentation Settings (ASA devices 9.6(1) or later)	
Enable Fragmentation before Encryption	Whether to enable fragmentation of IKEv2 messages. The Internet Key Exchange Version 2 (IKEv2) fragmentation protocol splits large IKEv2 message into a set of smaller ones, called IKE Fragment Messages. Fragmentation is supported on the ASA devices running the software version 9.6(1) or later.
Local MTU Size (ASA)	Enter the MTU size value. MTU size is used to divide the clear text packet into chunks. The MTU value used includes the IP header plus UDP header size. The default MTU size is 576.

Table 26-7 VPN Global Settings Page, IKEv2 Settings Tab (continued)

Element	Description
Fragmentation Mode (ASA)	Select one of the following: • CSCO—refers to the current Cisco Proprietary Fragmentation method. • IETF—refers to the method defined by the IETF standard: draft-ietf-ipsecme-ikev2-fragmentation. By default IETF is selected.

Understanding NAT in VPNs

Network Address Translation (NAT) enables devices that use internal IP addresses to send and receive data through the Internet. It converts private, internal LAN addresses into globally routable IP addresses when they try to access data on the Internet. In this way, NAT enables a small number of public IP addresses to provide global connectivity for a large number of hosts.

NAT enhances the stability of your hub-and-spoke VPN tunnels or remote access connections because resources required for VPN connections are not used for other purposes, and the VPN tunnel is kept available for traffic requiring complete security. Sites inside the VPN can use NAT through a split tunnel to exchange non secure traffic with outside devices, and they do not squander VPN bandwidth or overwhelm the hub at the tunnel head-end by directing nonessential traffic through it.

Security Manager supports NAT with dynamic IP addressing only, and applies to it an overload feature that permits what is known as port-level NAT or Port Address Translation (PAT). PAT uses port addressing to associate thousands of private NAT addresses with a small group of public IP address. PAT is used if the addressing requirements of your network exceed the available addresses in your dynamic NAT pool.

Note When you enable PAT on Cisco IOS routers, an additional NAT rule is implicitly created for split-tunneled traffic on deployment. This NAT rule, which denies VPN-tunneled traffic and permits all other traffic (using the external interface as the IP address pool), is not reflected as a router platform policy. You can remove the NAT rule by disabling this feature. For more information, see NAT Page: Dynamic Rules, page 24-10.

You can configure traffic to bypass NAT configuration on site-to-site VPN traffic. To bypass NAT configuration on Cisco IOS routers, make sure the **Do Not Translate VPN Traffic** option is selected in the NAT Dynamic Rule platform policy (see NAT Dynamic Rule Dialog Box, page 24-11). To exclude NAT on PIX Firewalls or ASA devices, make sure this option is selected in the NAT Translation Options platform policy (see Translation Options Page, page 24-17).

About NAT Traversal

NAT traversal is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow.

If the IP address of the VPN interface on the spoke is not globally routable, the NAT on the middle device replaces it with a new globally routable IP address. This change is made in the IPsec header, and violates the checksum of the spoke causing a mismatch with the hub's checksum calculation. This results in loss of connectivity between the hub and the spoke.

With NAT traversal, the spoke adds a UDP header to the payload. The NAT on the middle device changes the IP address in the UDP header, leaving the IPsec header and checksum intact. On a middle device that uses static NAT, you must provide the static NAT IP address (globally routable) on the inside interface. The static NAT IP address is provided for all traffic through that interface that requires NAT. However, if the middle device uses dynamic NAT where the NAT IP address is unknown, you must define dynamic crypto on the hub to serve any connection request from the spoke. Security Manager generates the required tunnel configuration for the spoke.

Note NAT traversal is enabled by default on routers running IOS version 12.3T and later. If you want to disable the NAT traversal feature, you must do this manually on the device or using a FlexConfig (see Chapter 7, "Managing FlexConfigs").

You can define global NAT settings on the NAT Settings tab of the Global VPN Settings page as described in Configuring VPN Global NAT Settings, page 26-42.

Configuring VPN Global NAT Settings

Use the NAT Settings tab of the Global Settings page to define global Network Address Translation (NAT) settings that enable devices that use internal IP addresses to send and receive data through the Internet.

Note For site-to-site VPNs, if you want to bypass NAT configuration on IOS routers, make sure that the **Do Not Translate VPN Traffic** option is selected in the NAT Dynamic Rule platform policy (see NAT Dynamic Rule Dialog Box, page 24-11). To exclude NAT on PIX Firewalls or ASA devices, make sure this option is selected in the NAT Translation Options platform policy (see Translation Options Page, page 24-17).

Navigation Path

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector. Click the **NAT Settings** tab.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one, then click the **NAT Settings** tab.
- For site-to-site VPNs, do one of the following:
  - Open the Site-to-Site VPN Manager Window, page 25-18, select a topology in the VPNs selector, then select **VPN Global Settings** in the Policies selector. Click the **NAT Settings** tab.
  - (Policy view) Select **Site-to-Site VPN > VPN Global Settings** from the Policy Types selector. Select an existing shared policy or create a new one, then click the **NAT Settings** tab.

Related Topics

- Understanding NAT in VPNs, page 26-41
- Configuring VPN Global Settings, page 26-30

Field Reference

Table 26-8 VPN Global Settings Page, NAT Settings Tab

Element	Description
Enable Traversal Keepalive Interval	Whether to enable NAT traversal keepalive. NAT traversal keepalive is used for the transmission of keepalive messages when there is a device (middle device) located between a VPN-connected hub and spoke, and that device performs NAT on the IPsec flow. If you select this option, configure the interval, in seconds, between the keepalive signals sent between the spoke and the middle device to indicate that the session is active. The value can be from 5 to 3600 seconds. The default is 10 seconds. **Note** On Cisco IOS routers, NAT traversal is enabled by default. If you want to disable the NAT traversal feature, you must do this manually on the device or by using a FlexConfig (see Chapter 7, "Managing FlexConfigs").
Enable Traversal over TCP TCP Ports (Remote access VPNs only.)	Supported on ASA and PIX 7.0+ devices. When selected, encapsulates both the IKE and IPsec protocols within a TCP packet and enables secure tunneling through both NAT and PAT devices and firewalls. If you select this option, specify the TCP ports for which you want to enable NAT traversal (NAT-T). You must configure TCP ports on the remote clients and on the VPN device. The client configuration must include at least one of the ports you set for the security appliance. You can enter up to 10 ports. **Tip** These ports are used for IKEv1 connections only. IKEv2 uses ports 500 and 4500 for NAT-T. Ensure that any ports that you specify are opened in the access rules for the applicable interface.
Enable PAT (Port Address Translation) on Split Tunneling for Spokes (Site-to-site VPNs only.)	Supported on Cisco IOS routers and Catalyst 6500/7600 devices. When selected, enables Port Address Translation (PAT) to be used for split-tunneled traffic on spokes in your VPN topology. PAT can associate thousands of private NAT addresses with a small group of public IP address through the use of port addressing. PAT is used if the addressing requirements of your network exceed the available addresses in your dynamic NAT pool. **Note** When you select this option, Security Manager implicitly creates an additional NAT rule for split-tunneled traffic on deployment. This NAT rule, which denies VPN-tunneled traffic and permits all other traffic (using the external interface as the IP address pool), is not reflected as a router platform policy. For information on creating or editing a dynamic NAT rule as a router platform policy, see NAT Page: Dynamic Rules, page 24-10.

Configuring VPN Global General Settings

Use the General Settings tab of the VPN Global Settings page to define fragmentation settings including maximum transmission unit (MTU) handling parameters for site-to-site and remote access VPNs.

Fragmentation breaks a packet into smaller units when it is transmitted over a physical interface that cannot support the original size of the packet. Fragmentation minimizes packet loss in a VPN tunnel, because it enables transmission of secured packets that might otherwise be too large to transmit. This is particularly relevant when using GRE, because any packet of more than 1420 bytes will not have enough room in its header for the additional 80 bytes that the combined use of IPsec and GRE adds to the packet payload.

The maximum transmission unit (MTU) specifies the maximum packet size, in bytes, that an interface can handle. If a packet exceeds the MTU, it is fragmented, typically after encryption. If the DF (Do Not Fragment) bit is set, the packet is dropped. A DF bit is a bit within the IP header that indicates if a device can fragment a packet. You must specify whether the device can clear, set, or copy the DF bit from the encapsulated header.

Because reassembly of an encrypted packet is difficult, fragmentation can degrade network performance. To prevent network performance problems, you can select **Enable Fragmentation Before Encryption** so that fragmentation occurs before encryption.

Navigation Path

- For remote access VPNs, do one of the following:
  - (Device View) Select **Remote Access VPN > Global Settings** from the Policy selector. Click the **General Settings** tab.
  - (Policy View) Select **Remote Access VPN > Global Settings** from the Policy Type selector. Select an existing policy or create a new one, then click the **General Settings** tab.
- For site-to-site VPNs, do one of the following:
  - Open the Site-to-Site VPN Manager Window, page 25-18, select a topology in the VPNs selector, then select **VPN Global Settings** in the Policies selector. Click the **General Settings** tab.
  - (Policy view) Select **Site-to-Site VPN > VPN Global Settings** from the Policy Types selector. Select an existing shared policy or create a new one, then click the **General Settings** tab.

Related Topics

- Configuring VPN Global Settings, page 26-30

Field Reference

Table 26-9 VPN Global Settings Page, General Settings Tab

Element	Description
Fragmentation Settings	
Fragmentation Mode Local MTU Size	Supported on Cisco IOS routers and Catalyst 6500/7600 devices. Fragmentation minimizes packet loss in a VPN tunnel when packets are transmitted over a physical interface that cannot support the original size of the packet. Select the fragmentation mode: • **No Fragmentation**—Do not fragment before IPsec encapsulation. After encapsulation, the device fragments packets that exceed the MTU setting before transmitting them through the public interface. • **End to End MTU Discovery**—Use ICMP messages to determine the maximum MTU. Use this option with IPsec VPNs. End-to-end MTU discovery uses Internet Control Message Protocol (ICMP) messages to determine the maximum MTU that a host can use to send a packet through the VPN tunnel without causing fragmentation. The MTU setting for each link in a transmission path is checked to ensure that no transmitted packet exceeds the smallest MTU in that path. The discovered MTU is used to decide whether fragmentation is necessary. If ICMP is blocked, MTU discovery fails and packets are either lost (if the DF bit is set) or fragmented after encryption (if the DF bit is not set). **Note** (Site-to-site VPNs) For Catalyst 6500/7600 devices, end-to-end path MTU discovery is supported only on images 12.2(33)SRA, 12.2(33)SRB, 12.2(33)SXH, 12.2(33)SXI or above. • **Local MTU Handling**—Set the MTU locally on the devices. This option is typically used when ICMP is blocked or in site-to-site IPsec/GRE VPNs. If you select this option, specify the local MTU size, which can be between 68 and 65535 bytes depending on the VPN interface.
DF Bit	Supported on Cisco IOS routers, Catalyst 6500/7600 devices, PIX 7.0+ and ASA devices. A Do Not Fragment (DF) bit within an IP header determines whether a device is allowed to fragment a packet. Select how to handle the DF bit: • **Copy**—Copy the DF bit from the encapsulated header in the current packet to all the device's packets. If the packet's DF bit is set to fragment, all future packets are fragmented. This is the default option. • **Set**—Set the DF bit in the packet you are sending. A large packet that exceeds the MTU is dropped and an ICMP message is sent to the packet's initiator. • **Clear**—Fragment packets regardless of the original DF bit setting. If ICMP is blocked, MTU discovery fails and packets are fragmented only after encryption.

Table 26-9 VPN Global Settings Page, General Settings Tab (continued)

Element	Description
Enable Fragmentation Before Encryption	Supported on Cisco IOS routers, Catalyst 6500/7600 devices, PIX 7.0+ and ASA devices. When selected, enables fragmentation to occur before encryption if the expected packet size exceeds the MTU. Look ahead Fragmentation (LAF) is used before encryption takes place to calculate the packet size that would result after encryption, depending on the transform sets configured on the IPsec SA. If the packet size exceeds the specified MTU, the packet will be fragmented before encryption.
Enable Notification on Disconnection	Supported on ASA and PIX 7.0+ devices. When selected, enables the device to notify qualified peers of sessions that are about to be disconnected. The peer receiving the alert decodes the reason and displays it in the event log or in a pop-up window. This feature is disabled by default. IPsec sessions might be dropped for several reasons, such as a security appliance shutdown or reboot, session idle timeout, maximum connection time exceeded, or administrator cut-off.
Enable Split Tunneling (Site-to-site VPN only.)	When selected (the default), enables you to configure split tunneling in your site-to-site VPN topology. Split tunneling allows you to transmit both secured and unsecured traffic on the same interface. Split tunneling requires that you specify exactly which traffic will be secured and what the destination of that traffic is, so that only the specified traffic enters the IPsec tunnel, while the rest is transmitted unencrypted across the public network.
Enable Spoke-to-Spoke Connectivity through the Hub	Supported on ASA and PIX 7.0+ devices. When selected, enables direct communication between spokes in a hub-and-spoke VPN topology in which the hub is an ASA or PIX 7.0+ device.
Enable Default Route	Supported on Cisco IOS routers and Catalyst 6500/7600 devices. When selected, the device uses the configured external interface as the default outbound route for all incoming traffic.
Do not reboot until all the sessions are terminated (ASA)	Select this option if you want the ASA to postpone a scheduled reboot until all active sessions terminate. This feature is disabled by default. **Note** The **crypto isakmp reload-wait** command in ASA software is supported only in System context for ASA devices that are on multiple context mode. However, since System context is not supported on VPN configurations, Security Manager does not generate this command for devices in VPN configuration, that are on multiple context mode. You must use the FlexConfig policies in System context for the **crypto isakmp reload-wait** command to work on devices that are on multiple context mode. FlexConfig policies allow you to configure device commands that are not otherwise supported by Security Manager. For more information, see Chapter 7, "Managing FlexConfigs".

Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs

If you want to use preshared key as your authentication method for IKEv1 negotiations, you must define a shared key for each tunnel between two peers that will be their shared secret for authenticating the connection. The key will be configured on each peer. If the key is not the same on both peers of the tunnel, the connection cannot be established. The peer addresses that are required for configuring the preshared key are deduced from the VPN topology.

Tip You can also use preshared keys for IKEv2 negotiations, but the configuration is different from the one used for IKEv1, as are the rules and requirements. For information on configuring preshared keys for IKEv2 negotiations, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

Preshared keys are configured on spokes. In a hub-and-spoke VPN topology, Security Manager mirrors the spoke's preshared key and configures it on its assigned hub, so that the key on the spoke and hub are the same. In a point-to-point VPN topology, you must configure the same preshared key on both peers. In a full mesh VPN topology, any two devices that are connected must have the same preshared key.

In a preshared key policy, you can use a specific key, or you can use automatically generated keys for peers participating in each communication session. Using automatically generated keys (the default method) is preferred, because security can be compromised if all connections in a VPN use the same preshared key.

Note Beginning with 4.16, Cisco Security Manager does not support IKEv1 related configuration for Firepower 9300 devices with distributed mode.

While discovering a VPN topology, where one of the devices is in cluster distributed mode (IKEv2 configured), and other is a non-cluster mode (IKEv1 and IKEv2 configured), Cisco Security Manager does not display any error. However, during preview config, the activity validation error is displayed to remove IKEv1 related configuration.

There are three methods for negotiating key information and setting up IKE security associations (SAs):

- Main mode address—Negotiation is based on IP address. Main mode provides the highest security because it has three two-way exchanges between the initiator and receiver. This is the default negotiation method.

  This method has three options for creating keys:

  - You can create a key for each peer, based on the unique IP address of each peer, providing high security.
  - You can create a group preshared key on a hub in a hub-and-spoke VPN topology, to be used for communication with any device in a specified subnet. Each peer is identified by its subnet, even if the IP address of the device is unknown. In a point-to-point or full mesh VPN topology, a group preshared key is created on the peers.
  - You can create a wildcard key on a hub in a hub-and-spoke VPN topology, or on a group containing hubs, to be used for dynamic crypto where a spoke does not have a fixed IP address or belong to a specific subnet. All spokes connecting to the hub have the same preshared key, which could compromise security. In a point-to-point or full mesh VPN topology, a wildcard key is created on the peers.

Note If you are configuring DMVPN with direct spoke-to-spoke connectivity, you create a wildcard key on the spokes.

- Main mode fully qualified domain name (FQDN)—Negotiation is based on DNS resolution, with no reliance on IP address. This option can only be used if the DNS resolution service is available for the host. It is useful when managing devices with dynamic IP addresses that have DNS resolution capabilities.
- Aggressive mode—Negotiation is based on hostname (without DNS resolution) and domain name. Aggressive mode is less secure than main mode. However, it provides more security than using group preshared keys if the IP address of the VPN interface on the host is unknown, and the FQDN of the dynamic IP peer is not DNS resolvable. This negotiation method is recommended for use with a GRE Dynamic IP or DMVPN failover and routing policy.

Related Topics

- Deciding Which Authentication Method to Use, page 26-8
- Configuring IKEv1 Preshared Key Policies, page 26-48

Configuring IKEv1 Preshared Key Policies

Use the IKEv1 Preshared Key page to define the preshared key configuration when using IKEv1 in a site-to-site VPN topology. For information on configuring preshared keys when using IKEv2, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

Note The preshared key policy does not apply to Easy VPN topologies.

Note Beginning with 4.16, Cisco Security Manager does not support IKEv1 preshared key configuration for Firepower 9300 devices with distributed mode.

To open the IKEv1 Preshared Key page:

- (Site-to-Site VPN Manager Window, page 25-18) Select a topology in the VPNs selector, then select **IKEv1 Preshared Key** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > IKEv1 Preshared Key** from the Policy Types selector. Select an existing shared policy or create a new one.

The following table explains the settings you can configure in this policy.

Table 26-10 IKEv1 Preshared Key Page

Element	Description
Key Specification	
Select whether to manually define the key (User Defined) or to have the key automatically generated. There are additional options you can configure when using auto generated keys.	

Table 26-10 IKEv1 Preshared Key Page (continued)

Element	Description
User Defined	When selected, enables you to use a manually defined preshared key. Enter the required preshared key in the **Key** field, then enter it again in the **Confirm** field.
Auto Generated	When selected, allocates a random key to the participating peers. This ensures security because a different key is generated for every hub-spoke connection. Auto Generated is the default selection. Auto generated is not a useful option when you do not manage all nodes in the VPN, for example, in the case of an Extranet VPN. **Note** The key is allocated during the first deployment to the devices and is used in all subsequent deployments to the same devices, until you select the Regenerate Key (Only in Next Deployment) check box.
Key Length	The required length of the preshared key to be automatically generated, from 1 to 127. The default is 24.
Same Key for All Tunnels	Unavailable in a point-to-point VPN topology. When selected, enables you to use the same auto-generated key for all tunnels. **Note** If you do not select this option, different keys are used for the tunnels, except in cases, such as DMVPN configuration, when different multipoint GRE interfaces in the same network must use the same preshared key.
Regenerate Key (Only in Next Deployment)	When selected, enables Security Manager to generate a new key for the next deployment to the devices. This is useful if it is possible that the secrecy of the keys might be compromised. When you submit the job for deployment, this check box is cleared. It does not remain selected because the new key will only be generated for the upcoming deployment, and not for subsequent deployments (unless you select it again).
Negotiation Method Select the type of negotiation method. The methods are explained in more detail in Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47.	

Table 26-10 IKEv1 Preshared Key Page (continued)

Element	Description
Main Mode Address	Use this negotiation method for exchanging key information if the IP address of the devices is known. Negotiation is based on IP address. Main mode provides the highest security because it has three two-way exchanges between the initiator and receiver. Main mode address is the default negotiation method. Select one of the following options to define the negotiation address type: • **Peer Address**—Negotiation is based on the unique IP address of each peer. A key is created for each peer, providing high security. This is the default. • **Subnet**—Creates a group preshared key on a hub in a hub-and-spoke topology to use for communication with any device in a specified subnet, even if the IP address of the device is unknown. Each peer is identified by its subnet. In a point-to-point or full mesh VPN topology, a group preshared key is created on the peers. Enter the subnet in the field provided, for example, 10.10.10.0/24. • **Wildcard**—Creates a wildcard key on a hub or on a group of hubs in a hub-and-spoke topology to use when a spoke does not have a fixed IP address or belong to a specific subnet. In this case, all spokes connecting to the hub have the same preshared key, which could compromise security. Use this option if a spoke in your hub-and-spoke VPN topology has a dynamic IP address. In a point-to-point or full mesh VPN topology, a wildcard key is created on the peers. **Note** When configuring DMVPN with direct spoke-to-spoke connectivity, you create a wildcard key on the spokes.
Main Mode FQDN	Select this negotiation method for exchanging key information if the IP address is not known and DNS resolution is available for the devices. Negotiation is based on DNS resolution, with no reliance on IP address.
Aggressive Mode	Available only in a hub-and-spoke VPN topology. Select this negotiation method for exchanging key information if the IP address is not known and DNS resolution might not be available on the devices. Negotiation is based on hostname and domain name. **Note** If direct spoke to spoke tunneling is enabled, you cannot use aggressive mode.

Related Topics

- Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47

Understanding Public Key Infrastructure Policies

Security Manager supports IPsec configuration with Certification Authority (CA) servers that manage certificate requests and issue certificates to devices in your VPN topology. You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and manage keys and certificates, providing centralized key management for the participating devices.

CA servers, also known as trustpoints, manage public CA certificate requests and issue certificates to participating IPsec network devices. When you use Certificates as the authentication method for IKE and IPsec proposal policies, peers are configured to obtain digital certificates from a CA server. With a CA server, you do not have to configure keys between all the encrypting devices. Instead, you individually enroll each participating device with the CA server, which is explicitly trusted to validate identities and create a digital certificate for the device. When this has been accomplished, each participating peer can validate the identities of the other participating peers and establish encrypted sessions with the public keys contained in the certificates.

CAs can also revoke certificates for peers that no longer participate in an IPsec VPN topology. Revoked certificates are either managed by an Online Certificate Status Protocol (OCSP) server or are listed in a certificate revocation list (CRL) stored on an LDAP server, which each peer can check before accepting a certificate from another peer.

PKI enrollment can be set up in a hierarchical framework consisting of multiple CAs. At the top of the hierarchy is a root CA, which holds a self-signed certificate. The trust within the entire hierarchy is derived from the RSA key pair of the root CA. Subordinate CAs within the hierarchy can enroll with either the root CA or with another subordinate CA. Within a hierarchical PKI, all enrolled peers can validate each other's certificate if the peers share a trusted root CA certificate or a common subordinate CA.

Keep the following in mind:

- PKI policies can be configured on Cisco IOS routers running version 12.3(7)T and later, PIX Firewalls, and Adaptive Security Appliance (ASA) devices for site-to-site and remote access VPNs.
- In site-to-site VPNs, you use the IKEv1 Public Key Infrastructure policy to identify CA servers for IKEv1 negotiations only. For IKEv2 negotiations, you identify the CA servers in the IKEv2 Authentication policy as described in Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.
- To save the RSA key pairs and the CA certificates between reloads permanently to Flash memory on a PIX Firewall release 6.3, you must configure the **ca save all** command. You can do this manually on the device or using a FlexConfig (see Chapter 7, "Managing FlexConfigs").

CA Server Authentication Methods

You can authenticate the CA server using one of the following methods:

- Using the Simple Certificate Enrollment Protocol (SCEP) to retrieve the CA's certificates from the CA server. Using SCEP, you establish a direct connection between your device and the CA server. Be sure your device is connected to the CA server before beginning the enrollment process. Because this method of retrieving CA certificates for routers is interactive, you can deploy your PKI policy to live devices only, not to files.

Note When using SCEP, you must enter the fingerprint for the CA server. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. You can obtain the CA's fingerprint by contacting the server directly, or by entering the following address in a web browser: **http://<URLHostName>/certsrv/mscep/mscep.dll**.

- Manually creating an enrollment request that you can submit to a CA server offline, by copying the CA server's certificates from another device.

  Use this method if your device cannot establish a direct connection to the CA server or if you want to generate an enrollment request and send it to the server at a later time.

  Note This method enables you to deploy the PKI policy either to devices or to files.

  For more information, see PKI Enrollment Dialog Box, page 26-58.

Note You can also use Cisco Secure Device Provisioning (SDP) to enroll for a certificate for a router. For more information about using SDP for certificate enrollment, see Secure Device Provisioning on Cisco IOS Routers, page 63-81.

The following topics explain Public Key Infrastructure configuration in more detail:

- Requirements for Successful PKI Enrollment, page 26-52
- Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54
- Defining Multiple IKEv1 CA Servers for Site-to-Site VPNs, page 26-55
- Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56
- PKI Enrollment Dialog Box, page 26-58

Requirements for Successful PKI Enrollment

The following are prerequisites for configuring a PKI policy in your network:

- For IKEv1, the IKE proposal must specify Certificate for the IKE authentication method. See Configuring IKEv1 Proposal Policy Objects, page 26-10.
- The domain name must be defined on the devices for PKI enrollment to be successful (unless you specify the CA server nickname).
- To enroll with the CA server directly, you must specify the server's enrollment URL.
- To enroll with the CA server by means of a TFTP server, you must ensure that the CA certificates file is saved to the TFTP server. After deployment of the PKI policy, you must copy the certificate request from your TFTP server to the CA server.
- You may specify an RSA public key to use in the enrollment request. If you do not specify an RSA key pair, the Fully Qualified domain Name (FQDN) key will be used.

  If using RSA keys, once the certificate has been granted, the public key is included in the certificate so that peers can use it to encrypt data sent to the device. The private key is kept on the device and used to decrypt data sent by peers, and to digitally sign transactions when negotiating with peers. You can use an existing key pair or generate a new one. If you want to generate a new key pair to use in the certificate for router devices, you must also specify the modulus to determine the size of the key.

  For more information, see PKI Enrollment Dialog Box—Enrollment Parameters Tab, page 26-63.
- If you are making a PKI enrollment request on a Cisco Easy VPN IPsec remote access system, you must configure each remote component (spoke) with the name of the user group to which it connects. You specify this information in the Organization Unit (OU) field in the Certificate Subject Name tab of the PKI Enrollment Editor dialog box.

Note You do not need to configure the name of the user group on the hub (Easy VPN server).

For more information, see PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66.

- To deploy PKI policies to files (not to live devices), the following prerequisites must be met:
  - Routers must run Cisco IOS Software 12.3(7)T or later.
  - CA authentication certificates must be cut and pasted into the Security Manager user interface (so that CA authentication is not interactive and does not require communication with the live device).
- If you are deploying to live devices, the PKI server must be online.
- Security Manager supports the Microsoft, Verisign, and Entrust PKIs.
- Security Manager supports Cisco IOS Certificate Servers. The Cisco IOS Certificate Server feature embeds a simple certificate server, with limited CA functionality, into the Cisco IOS software. An IOS Certificate Server can be configured as a FlexConfig policy. For more information, see Chapter 7, "Managing FlexConfigs".
- To configure PKI with AAA authorization that uses the entire subject name on an IOS router, use the predefined FlexConfig object named IOS_PKI_WITH_AAA.

Prerequisites for PKI Enrollment Using TFTP

If you do not have constant direct access to the CA server, you can enroll using TFTP if your devices are routers running Cisco IOS Software 12.3(7)T or later.

On deployment, Security Manager generates the corresponding CA trustpoint command and authenticate command. The trustpoint command is configured with the enrollment URL tftp://<certserver> <file_specification> entry to retrieve the CA certificate using TFTP. If file_specification is not specified, the FQDN of the router is used.

Before using this option, you must make sure that the CA certificates file (.ca) is saved on the TFTP server. To do this, use this procedure:

1. Connect to http://servername/certsrv, where servername is the name of the Windows 2000 web server on which the CA you want to access is located.
2. Select **Retrieve the CA certificate or certificate revocation list**, then click **Next**.
3. Select **Base64 encoded**, then click **Download CA certificate**.
4. Save the .crt file as a .ca file on the TFTP server using your browser's Save As function.

After deployment, you must transfer the certificate request generated by Security Manager on the TFTP server to the CA, and then transfer the device's certificates from the CA to the device.

Transferring the Certificate Request from the TFTP Server to the CA Server

Security Manager creates a PKCS#10 formatted enrollment request (.req) on the TFTP server. You must transfer it to the PKI server using this procedure:

1. Connect to http://servername/certsrv, where servername is the name of the Windows 2000 web server where the CA you want to access is located.
2. Select **Request a certificate**, then click **Next**.
3. Select **Advanced request**, then click **Next**.
4. Select **Submit a certificate request using a base64 encoded PKCS #10 file or a renewal request using a base64 encoded PKCS #7 file**, then click Next.

5. Either select browse for a file (and browse to the TFTP server and select the .req file) or open the just received by TFTP .req file with WordPad/Notepad and copy/paste the contents in the first window.
6. Export the .crt file from the CA and put it on the TFTP server.
7. Configure the 'crypto ca import <label> certificate' to import the device's certificates from the tftp server.

Related Topics

- Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54
- Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56
- PKI Enrollment Dialog Box, page 26-58
- Configuring a User Group Policy for Easy VPN, page 28-14

Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs

You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and to manage keys and certificates. Certification Authority (CA) servers are used to manage these certificate requests and issue certificates to the participating devices in your VPN topology.

In Security Manager, CA servers are predefined as PKI enrollment objects that you can use in your PKI policies. A PKI enrollment object contains the server information and enrollment parameters that are required for creating enrollment requests for CA certificates.

For more information about Public Key Infrastructure policies, see Understanding Public Key Infrastructure Policies, page 26-51.

This procedure describes how to specify the CA server that will be used to create an IKEv1 Public Key Infrastructure (PKI) policy in your VPN topology.

Tip For information on specifying CA servers for use in IKEv2 negotiations, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

Before You Begin

For important information about successfully configuring PKI, see Requirements for Successful PKI Enrollment, page 26-52.

Related Topics

- Defining Multiple IKEv1 CA Servers for Site-to-Site VPNs, page 26-55
- Deciding Which Authentication Method to Use, page 26-8
- Filtering Items in Selectors, page 1-45

Step 1 Do one of the following:

- (Site-to-Site VPN Manager Window) Select an existing topology and then select **IKEv1 Public Key Infrastructure** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > IKEv1 Public Key Infrastructure**, and then select an existing policy or create a new one.

The Public Key Infrastructure page opens, displaying the currently selected CA server, if any, in the **Selected** field.

Step 2 Select the PKI enrollment policy object that defines the desired CA server in the Available CA Servers list. You can do the following to modify the listed objects:

- To add a new PKI enrollment object, click the **Create (+)** button. The Add PKI Enrollment dialog box opens. For detailed information about the attributes of a PKI enrollment object, see PKI Enrollment Dialog Box, page 26-58.
- To change the configuration of an existing object, select it and click the **Edit (pencil)** button.

Note If you are making a PKI enrollment request on an Easy VPN topology, you must configure each remote component (spoke) with the name of the user group to which it connects. You specify this information in the Organization Unit (OU) field in the Certificate Subject Name tab of the PKI Enrollment dialog box. You do not need to configure the name of the user group on the hub (Easy VPN server). For more information, see PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66.

Defining Multiple IKEv1 CA Servers for Site-to-Site VPNs

You can select only one CA server when defining an IKEv1 Public Key Infrastructure (PKI) policy on a site-to-site VPN. This creates a problem when the devices in the VPN enroll with different CA servers when using IKEv1. For example, the spoke devices might enroll with a different CA server than the hub, or the spokes in one part of the VPN might enroll with a different CA server than the spokes in another part of the VPN.

Tip When using IKEv2, you can configure different CA servers for various devices by creating overrides for the IKEv2 Authentication policy global settings rather than creating device-level overrides for PKI enrollment policy objects. However, you can also use device-level overrides for IKEv2 as described in this topic. For information on configuring CA servers for IKEv2, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

To define an IKEv1 PKI policy, you select a PKI enrollment object that specifies the CA server to which the devices should enroll. Although by default the policy object refers globally to a single CA server, you can use device-level overrides to have the object refer to a different CA server on selected devices.

For example, if PKI enrollment object PKI_1 refers to a CA server named CA_1, you can create a device-level override for selected devices that has PKI_1 refer to a different CA server, for example, CA_2. Theoretically, you can use overrides to define a different CA server for each device in the VPN.

This procedure describes the basic steps for creating overrides for PKI enrollment objects.

Note You can also use device-level overrides when the CA servers are arranged in a PKI hierarchy beneath a common, trusted CA server. To do this, you must ensure that both the global definition of the object and the device-level override specify the trusted CA server in the Trusted CA Hierarchy tab of the PKI Enrollment dialog box. See PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67.

Related Topics

- Understanding Public Key Infrastructure Policies, page 26-51
- Deciding Which Authentication Method to Use, page 26-8

Step 1 To create the PKI enrollment object, open the PKI Enrollment dialog box. You can access this dialog box in two ways:

- From the Public Key Infrastructure policy—Click the **Create (+)** button beneath the Selected field. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.
- From the Policy Object Manager (select **Manage > Policy Objects**)—Select PKI Enrollments from the Object Type selector, then click the **New Object (+)** button.

Step 2 Define the global definition of the PKI enrollment object, including the CA server to which the object refers. Be sure to select **Allow Value Override per Device**. This option makes the object overridable on individual devices. See PKI Enrollment Dialog Box, page 26-58.

Base the global definition of the object on the CA server that is used by the most devices in the VPN. Doing this reduces the number of device-level overrides that are required.

Step 3 When you finish defining the PKI enrollment object, click **OK**. As a result:

- If you accessed the dialog box through the PKI policy, the new object appears in the Selected field of the policy page.
- If you accessed the dialog box using the Policy Object Manager, the new object appears in the work area of the Policy Object Manager window. A green check mark in the Overridable column indicates that device-level overrides can be created for this object. (The check mark does not indicate whether any overrides actually exist.)

Step 4 Create the device-level overrides for the PKI enrollment object. You can do this in one of two ways:

- From Device Properties (with the device selected in Device view, select **Tools > Device Properties**)—This option is recommended when you want to create a device-level override for a single device. In the device properties, select **Policy Object Overrides > PKI Enrollments**, select the PKI enrollment object that you want to override, then click the **Create Override** button. You can then define the content of the override, including the CA server defined by the object.

  For more information, see Creating or Editing Object Overrides for a Single Device, page 6-19.
- From the Policy Object Manager—This option is recommended when you want to create a device-level override for multiple devices at the same time. Double-click the green check mark in the Overridable column, select the devices to which the override should apply, then define the content of the override, including the CA server defined by the object.

  For more information, see Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19.

Configuring Public Key Infrastructure Policies for Remote Access VPNs

You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and to manage keys and certificates. Certification Authority (CA) servers are used to manage these certificate requests and issue certificates to users who connect to your IPsec or SSL remote access VPN.

In Security Manager, CA servers are predefined as PKI enrollment objects that you can use in your PKI policies. A PKI enrollment object contains the server information and enrollment parameters that are required for creating enrollment requests for CA certificates.

For more information about Public Key Infrastructure policies, see Understanding Public Key Infrastructure Policies, page 26-51.

Note Beginning with version 4.12, Security Manager provides support for Public Key Infrastructure policy for ASA multi-context devices running the software version 9.5(2) or later.

This procedure describes how to specify the CA servers that will be used to create a Public Key Infrastructure (PKI) policy in your remote access VPN.

Before You Begin

Keep the following in mind:

- For important information about successfully configuring PKI, see Requirements for Successful PKI Enrollment, page 26-52.
- The **IKE Proposal** policy for IPsec remote access VPNs should use an IKE Proposal object that requires certificate authorization when configuring IKEv1.

- For remote access VPNs defined on an ASA or PIX 7.x+ device, be aware that the Public Key Infrastructure policy is directly related to the following policies. Any trustpoints defined in these policies must also be selected in the Public Key Infrastructure policy; they are not automatically added to the policy. You might want to first configure these policies to determine which PKI enrollment objects are required in your remote access VPNs.
  - **Connection Profiles**—When you create a IPsec connection profile for which CA trustpoints should be used, you select the PKI enrollment object that identifies the trustpoint on the IPsec tab.
  - **SSL VPN Access**—You can configure trustpoints for each interface and also a fallback trustpoint.
  - **Global Settings, IKEv2 Settings tab**—For IKEv2 IPsec, you must specify a global trustpoint.

Related Topics

- Deciding Which Authentication Method to Use, page 26-8
- Filtering Items in Selectors, page 1-45

Step 1 Do one of the following:

- (Device view) Select **Remote Access VPN > Public Key Infrastructure** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Public Key Infrastructure** from the Policy Type selector. Select an existing policy or create a new one.

The Public Key Infrastructure page opens, displaying the currently available and selected CA servers (PKI enrollment objects), if any.

Step 2 Select the PKI enrollment policy objects that define the desired CA servers in the Available CA Servers list and click **>>** to move them to the Selected CA Servers list. You can remove undesired objects by selecting them in the selected list and clicking **<<**.

Note When configuring IKEv2 in a Site-to-Site VPN, and choosing PKI as the authentication method, you are required to specify the object name that must be listed here, under Selected CA Servers (see Step 2 of Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68). Hence, ensure that you include the required CA Servers in the Selected CA Servers list.

For ASA and PIX 7.x+ devices, the list of selected PKI enrollment objects must include all objects that are specified in the connection profiles defined for the remote access VPN. For more information on connection profiles, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7. Also, any trustpoints configured for IKEv2 on the Global Settings policy must be included; see Configuring VPN Global IKEv2 Settings, page 26-37.

You can do the following to modify the listed objects:

- To add a new PKI enrollment object, click the **Create (+)** button below the list of available servers. The Add PKI Enrollment dialog box opens. For detailed information about the attributes of a PKI enrollment object, see PKI Enrollment Dialog Box, page 26-58.
- To change the configuration of an existing object, select it in either list and click the **Edit (pencil)** button.

PKI Enrollment Dialog Box

Use the PKI Enrollment dialog box to view, create, copy, or edit Public-Key Infrastructure (PKI) enrollment objects. A PKI enrollment object represents an external certification authority (CA) server that responds to certificate requests from devices in the network.

You can create PKI enrollment objects to define the properties of a CA server used when devices exchange certificates as part of an IPsec network. When you create a PKI enrollment object, you define a name for the server and the URL for enrollment. You must specify whether the devices you wish to enroll with this server should retrieve the CA server's own certificate using the Simple Certificate Enrollment Process (SCEP) or use a certificate that you have entered manually into the device configuration. You must also select the method of support used by the CA server for revocation checking.

Note You do not have to define enrollment parameters in order to create or import a trustpoint in Security Manager.

In addition, you can optionally define the following:

- Whether the CA server is acting as a Registration Authority (RA) server.
- Enrollment parameters, including retry settings and RSA key pair settings.
- Additional attributes to include in the certificate request.
- The list of trusted CA servers located above this server in the PKI hierarchy.

Navigation Path

Select **Manage > Policy Objects**, then select **PKI Enrollments** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Tip You can also open this dialog box from the **Public Key Infrastructure** policy for remote access or site-to-site VPNs.

Related Topics

- Understanding Public Key Infrastructure Policies, page 26-51
- Requirements for Successful PKI Enrollment, page 26-52
- Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54
- Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68
- Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56
- Policy Object Manager, page 6-4

Field Reference

Table 26-11 PKI Enrollment Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
CA Information tab	Use this tab to enter settings related to the Certificate Authority server, its certificate, and its level of revocation checking support. For information on the specific settings, see PKI Enrollment Dialog Box—CA Information Tab, page 26-60.
Enrollment Parameters tab	Use this tab to enter settings related to PKI enrollment. For information on the specific settings, see PKI Enrollment Dialog Box—Enrollment Parameters Tab, page 26-63. **Note** You do not have to define enrollment parameters in order to create or import a trustpoint in Security Manager.
Certificate Subject Name tab	Use this tab to enter optional information to be included in the certificate, including subject attributes. For information on the specific settings, see PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66.
Trusted CA Hierarchy tab	Use this tab to define trusted CA servers that are arranged in a hierarchical framework. For information on the specific settings, see PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

PKI Enrollment Dialog Box—CA Information Tab

Use the CA Information tab of the PKI Enrollment dialog box to:

- Define the name and location of the external certificate authority (CA) server.
- Manually paste the certificate, if known.
- Define the server's level of support for revocation checking.

Navigation Path

Go to the PKI Enrollment dialog box and click the **CA Information** tab. For information on opening the dialog box, see PKI Enrollment Dialog Box, page 26-58.

Related Topics

- PKI Enrollment Dialog Box—Enrollment Parameters Tab, page 26-63
- PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66
- PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67

Field Reference

Table 26-12 PKI Enrollment Dialog Box—CA Information Tab

Element	Description
CA Server Nickname	The name used to identify the CA server in the certificate request. If you leave this field blank, the domain name is used. You must leave this field blank for Verisign CAs. Also, keep the following in mind: • You cannot configure two CA servers with the same name but different URLs on the same device. • The CA name cannot match the name of a trusted CA configured as part of the same PKI enrollment object (as defined on the PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67). • When the device is configured as part of a VPN, do not configure a device-level override that uses the same CA name as that of the CA server used by any of the peers. (This is not a problem when the device and its peers use a tiered PKI hierarchy.)
Enrollment Type	The type of enrollment you want to perform. Security Manager completes the enrollment only if you configure URL enrollment. If you select another type, you must complete the enrollment using your own methods. • Self-Signed Certificate (ASA only)—To configure the **enrollment self** command. • Terminal (ASA only)—To configure the **enrollment terminal** command. • URL—To configure the URL for the CA server so that you can complete automatic enrollment. • None—Do not configure any enrollment command.
Protocols	Specify whether you want to configure an SCEP CA URL or a CMP CA URL.

Table 26-12 PKI Enrollment Dialog Box—CA Information Tab (continued)

Element	Description
Enrollment URL (URL enrollment only.)	The URL of the CA server to which devices should attempt to enroll. The URL can be in the following formats: • SCEP—Uses an HTTP URL in the form of **http://CA_name:port**, where CA_name is the host DNS name or IP address of the CA server. The port number is mandatory. • TFTP—Uses the format **tftp://certserver/file_specification**. Use this option when you do not have direct access to the CA server. The TFTP server transfers certificate requests and certificates. • Other supported formats include: bootflash, cns, flash, ftp, null, nvram, rcp, scp, system. **Note** If the CA cgi-bin script location at the CA is not the default (/cgi-bin/pkiclient.exe), you must also include the nonstandard script location in the URL, in the form of **http://CA_name:port/script_location**, where script_location is the full path to the CA scripts.
CA Certificate Source Fingerprint Certificate (URL enrollment only.)	How to obtain the certificate: • **Retrieve CA Certificate Using SCEP** (the default)—Have the router retrieve the certificate from the CA server using the Simple Certificate Enrollment Process (SCEP). Enter the fingerprint for the CA server in hexadecimal format. If the value you enter does not match the fingerprint on the certificate, the certificate is rejected. Using the **Fingerprint** to verify the authenticity of the CA's certificate helps prevent an unauthorized party from substituting a fake certificate in place of the real one. **Tip** You can obtain the CA's fingerprint by contacting the server directly, or by entering the following address in a web browser: **http://URLHostName/certsrv/mscep/mscep.dll**. Using the fingerprint is supported only on Cisco IOS software releases 12.3(12) or higher, 12.3(14)T or higher, 12.4 or higher (including 15.x), 12.2(33)XNA or higher. • **Enter CA Certificate from CA Server Manually**—Copy and Paste up to three certificates from another device into the **Certificate** field (using your browser's Paste function or the Ctrl-V keyboard shortcut). Use this option when you want the PKI enrollment object to represent predefined certificates. Each certificate must begin with the word "certificate" and end with the word "quit". CMP authentication requires base 64 encoded CA certificate for authentication. For CMP, we can configure base 64 encoded CA certificate in this field. Copy and paste the base 64 encoded CA certificate from CA server and end with the word "quit". **Note** Enter the certificate details within the words, ----BEGIN CERTIFICATE---- and ----END CERTIFICATE----.

Table 26-12 PKI Enrollment Dialog Box—CA Information Tab (continued)

Element	Description
CA Certificate Check	Certificates without the CA flag now cannot be installed on the ASA as CA certificates by default. The basic constraints extension identifies whether the subject of the certificate is a CA and the maximum depth of valid certification paths that include this certificate. Beginning with version 4.9, you can use Security Manager to configure the ASA to allow installation of these certificates if desired. This feature is supported only in devices running ASA software version 9.4(1) or higher. CA Certificate check is enabled by default.
Revocation Check Support	The type of certificate revocation checking to be performed: • Checking Not Performed—This is the default. The device does not perform any revocation checking, even if a CRL is on the device. • CRL Check Required—The device must check a CRL. If no CRL exists on the device and the device cannot obtain one, certificates are rejected and a tunnel cannot be established. • OCSP Check Required—The device must check revocation status from an OCSP server. If this check fails, certificates are rejected. • CRL Check Attempted—The device tries to download the latest CRL from the specified LDAP server. If the download fails, however, certificates are accepted. • OCSP Check Attempted—The device tries to check revocation status from an OCSP server. If this fails, however, certificates are accepted. • CRL or OCSP Check Required—The device first checks for a CRL. If a CRL does not exist or cannot be obtained, the device tries to check revocation status from an OCSP server. If both options fail, certificates are rejected. • OCSP or CRL Check Required—The device first tries to check revocation status from an OCSP server. If this fails, the device checks for a CRL. If both options fail, certificates are rejected. • CRL and OCSP Checks Attempted—The device first checks for a CRL. If a CRL does not exist or cannot be obtained, the device tries to check revocation status from an OCSP server. If both options fail, however, certificates are accepted. • OCSP and CRL Checks Attempted—The device first tries to check revocation status from an OCSP server. If this fails, the device tries to download the latest CRL. If both options fail, however, certificates are accepted.
OCSP Server URL	The URL of the OCSP server checking for revocation if you require OCSP checks. This URL must start with **http://**
CRL Server URL	The URL of the LDAP server from which the CRL can be downloaded if you require CRL checks. This URL must start with **ldap://** **Note** You must include a port number in the URL when using this AAA server on ASA devices, otherwise LDAP will fail.

Table 26-12 PKI Enrollment Dialog Box—CA Information Tab (continued)

Element	Description
Enable Registration Authority Mode (PIX 6.3)	For PIX 6.3 devices, whether the CA server operates in RA (Registration Authority) mode. A Registration Authority is a server that acts as a proxy for the actual CA so that CA operations can continue when the CA server is offline. **Note** Cisco IOS routers configure RA mode automatically, if required.

PKI Enrollment Dialog Box—Enrollment Parameters Tab

Use the Enrollment Parameters tab of the PKI Enrollment dialog box to define the retry settings to use when the device contacts the CA server as well as the settings for generating the RSA key pair to associate with the certificate.

If the PKI enrollment object represents a Microsoft CA, you can define the challenge password required to validate the router's identity.

Note You do not have to define enrollment parameters in order to create or import a trustpoint in Security Manager.

Navigation Path

Go to the PKI Enrollment dialog box and click the **Enrollment Parameters** tab. For information on opening the dialog box, see PKI Enrollment Dialog Box, page 26-58.

Related Topics

- PKI Enrollment Dialog Box—CA Information Tab, page 26-60
- PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66
- PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67

Field Reference

Table 26-13 PKI Enrollment Dialog Box—Enrollment Parameters Tab

Element	Description
Challenge Password Confirm	The password used by the CA server to validate the identity of the device. This password is mandatory for PIX 6.3 devices, but optional for PIX/ASA 7.0+ devices and Cisco IOS routers. You can obtain the password by contacting the CA server directly or by entering the following address in a web browser: **http://URLHostName/certsrv/mscep/mscep.dll**. The password is good for 60 minutes from the time you obtain it from the CA server. Therefore, it is important that you deploy the password as soon as possible after you create it. **Note** Each password is valid for a single enrollment by a single device. Therefore, we do not recommend that you assign a PKI enrollment object where this field is defined to a VPN, unless you first configure a device-level override for each device in the VPN. For more information, see Understanding Policy Object Overrides for Individual Devices, page 6-18.
Retry Period	The interval between certificate request attempts, in minutes. Values can be 1 to 60 minutes. The default is 1 minute.
Retry Count	The number of retries that should be made if no certificate is issued upon the first request. Values can be 1 to 100. The default is 10.
Certificate Auto-Enrollment (IOS devices only)	The percentage of the current certificate's lifetime after which the router requests a new certificate. For example, if you enter 70, the router requests a new certificate after 70% of the lifetime of the current certificate has been reached. Values range from 10% to 100%. If you do not specify a value, the router requests a new certificate after the old certificate expires.
Enable Auto-Enrollment	When enabled, certificates are automatically requested based on configurable triggers. The following specific parameters can also be further configured: • whether or not CMPv2 update will be used • when it will be triggered • whether the current key pair will be used or a new key pair will be generated
Certificate Auto-Enrollment (ASA 9.7.1 onwards)	The percentage of the current certificate's lifetime after which the router requests a new certificate. For example, if you enter 50, the router requests a new certificate after 50% of the lifetime of the current certificate has been reached. Values range from 10% to 99%. If you do not specify a value, the router requests a new certificate after the old certificate expires. **Note** The default value is 70%.
Auto Enroll Regenerate Key (ASA 9.7.1 onwards)	Select to generate a new key while renewing the certificate.

Table 26-13 PKI Enrollment Dialog Box—Enrollment Parameters Tab (continued)

Element	Description
Regenerate Key Pair (ASA 9.7.1 onwards)	Select to regenerate a new key pair before enrolling the trustpoint request.
Shared Key (ASA 9.7.1 onwards)	Specify the user credentials obtained from the CA, out of band. This will be used by the CA and ASA to confirm the authenticity and integrity of the messages that they exchange. The key length cannot exceed 64 characters. **Note** The shared key must be in the format, 'reference: *shared key*'.
Signing Certificate (ASA 9.7.1 onwards)	Specify the name of the trust point that contains a previously issued device certificate to be used to sign the CMP enrollment request.
Note For the CMP protocol, options like Certificate, Shared Key or Signing Certificate will not be discovered, for security reasons. As a result, the PKI enrollment dialog will be creating a override on rediscovery.	
Key Pair	New key pairs will be automatically generated for all CMP manual and automatic enrollments. To support this, we the ability to configure key pair parameters in the trust point has been added. Select the algorithm - RSA or EDCSA, that should be used to generate the key pair. **Note** The RSA algorithm has the following modulus options: 1024| 2048| 4096| 512|768 and the EDCSA algorithm, has the following elliptic-curve options 256|384|521 to generate a key pair.
Include Device's Serial Number	Whether to include the serial number of the device in the certificate. **Tip** The CA uses the serial number to either authenticate certificates or to later associate a certificate with a particular device. If you are in doubt, include the serial number, as it is useful for debugging purposes.
RSA Key Pair Name (PIX 7.0+, ASA, IOS devices only.)	If the key pair you want to associate with the certificate already exists, this field specifies the name of that key pair. If the key pair does not exist, this field specifies the name to assign to the key pair that will be generated during enrollment. **Note** If you do not specify an RSA key pair, the fully qualified domain name (FQDN) key pair is used instead. On PIX and ASA devices, the key pair must exist on the device before deployment.
RSA Key Size (IOS devices only.)	If the key pair does not exist, defines the desired key size (modulus), in bits. If you want a modulus between 512 and 1024, enter an integer that is a multiple of 64. If you want a value higher than 1024, enter 1536 or 2048. The recommended size is 1024. **Note** The larger the modulus size, the more secure the key. However, keys with larger modulus sizes take longer to generate (a minute or more when larger than 512 bits) and longer to process when exchanged.

Table 26-13 PKI Enrollment Dialog Box—Enrollment Parameters Tab (continued)

Element	Description
RSA Encryption Key Size (IOS devices only.)	The size of the second key, which is used to request separate encryption, signature keys, and certificates.
Source Interface (IOS devices and ASA 9.5(1) or higher)	The source address for all outgoing connections sent to a CA or LDAP server during authentication, enrollment, and when obtaining a revocation list. This parameter may be necessary when the CA server or LDAP server cannot respond to the address from which the connection originated (for example, due to a firewall). If you do not define a value in this field, the address of the outgoing interface is used. Enter the name of an interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. **Note** Security Manager 4.9 supports separate routing table for Management traffic on devices running ASA 9.5(1) or higher. This functionality enables to completely segregate management traffic from other data traffic on the ASA. Apart from IOS devices you can now select ASA devices running the software version 9.5(1) or higher.

PKI Enrollment Dialog Box—Certificate Subject Name Tab

Use the Certificate Subject Name tab of the PKI Enrollment dialog box to optionally define additional information about the device in certificate requests sent to the CA server. This information is placed in the certificate and can be viewed by any party who receives the certificate from the router.

Enter all information using the standard LDAP X.500 format.

Navigation Path

Go to the PKI Enrollment dialog box and click the **Certificate Subject Name** tab. For information on opening the dialog box, see PKI Enrollment Dialog Box, page 26-58.

Related Topics

- PKI Enrollment Dialog Box—CA Information Tab, page 26-60
- PKI Enrollment Dialog Box—Enrollment Parameters Tab, page 26-63
- PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab, page 26-67

Field Reference

Table 26-14 PKI Enrollment Dialog Box—Certificate Subject Name Tab

Element	Description
Include FQDN	Whether to include the device's fully qualified domain name (FQDN) in the certificate request. The name is taken from the Hostname policy (ensure that you specify both the hostname and domain name in the policy to get a valid fully-qualified domain name). If you do not configure the Hostname policy, the name is derived from the display name for the device in Security Manager, *display_name***.null**, which is unlikely to give you the desired results.
Include Device's IP Address	The interface whose IP address is included in the certificate request. Enter the name of the interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Common Name (CN)	The X.500 common name to include in the certificate.
Organization Unit (OU)	The name of the organization unit (for example, a department name) to include in the certificate. **Note** When you configure PKI enrollment objects for Cisco Easy VPN Remote components, this field must contain the name of the client group to which the component connects. Otherwise, the component will not be able to connect. Although this information is not required for the Easy VPN Server, including it does not create configuration problems. For more information about Easy VPN, see Understanding Easy VPN, page 28-1.
Organization (O)	The organization or company name to include in the certificate.
Locality (L)	The locality to include in the certificate.
State (ST)	The state or province to include in the certificate.
Country (C)	The country to include in the certificate.
Email (E)	The email address to include in the certificate.

PKI Enrollment Dialog Box—Trusted CA Hierarchy Tab

Use the Trusted CA Hierarchy tab of the PKI Enrollment dialog box to define the trusted CA servers within an hierarchical PKI framework. Within this framework, all enrolled peers can validate each other's certificates if they share a trusted root CA certificate or a common subordinate CA.

Select the CA servers (as defined as PKI enrollment objects) to include in the hierarchy in the Available Servers list and click >> to move them to the selected list. You can do the reverse to remove servers.

If the PKI enrollment object you need is not yet defined, click the **Create (+)** button beneath the available servers list to create the object. You can also select an object and click the **Edit** button to change its definition, if needed.

Navigation Path

Go to the PKI Enrollment dialog box and click the **Trusted CA Hierarchy** tab. For information on opening the dialog box, see PKI Enrollment Dialog Box, page 26-58.

Related Topics

- PKI Enrollment Dialog Box—CA Information Tab, page 26-60
- PKI Enrollment Dialog Box—Enrollment Parameters Tab, page 26-63
- PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66

Configuring IKEv2 Authentication in Site-to-Site VPNs

When you configure IKE version 2 (IKEv2) in a site-to-site VPN, you must configure the IKEv2 Authentication policy to define authentication settings. Unlike IKEv1, authentication settings are not part of the IKEv2 proposal.

In Security Manager, when you configure IKEv2 authentication for a site-to-site VPN, you configure default settings that will be used in the VPN topology. You can then configure exceptions to the default, specifying different preshared keys or trustpoints for specific segments of the VPN. You can use a mixture of preshared keys and trustpoints, for example, configuring a global preshared key, but trustpoints for selected members of the VPN.

Configuring asymmetric authentication for IKEv2 tunnels

IKEv2 allows you to use asymmetric authentication, unlike IKEv1. This means that two peers can have different preshared keys, different trustpoints, or one peer could use a preshared key and the other peer could use a trustpoint. In Security Manager, you can configure asymmetric authentication by doing any of the following:

- On the Global IKEv2 Authentication Settings tab, you can configure different preshared keys if you elect to auto-generate keys and **do not** select the Same Keys for All Tunnels or the Same Key at Tunnel Endpoints option. A different preshared key is generated for each end of each tunnel.
- On the Override IKEv2 Authentication Settings tab, you can create overrides for the global settings. You add overrides that specify different keys or trustpoints for subsets of local and remote peers. Because you can create more than one override for a device or a specific tunnel, you can configure a set of preshared keys and trustpoints from which peers will authenticate.

Tip The IKEv2 Authentication policy is not a shared policy. You must configure the policy for each VPN topology in which you support IKEv2 negotiations. You cannot configure global IKEv2 authentication options for use by all of your VPN topologies. When using the Create VPN wizard, even if you elect to support IKEv2, the IKEv2 Authentication policy is never configured.

Before You Begin

The IKEv2 Authentication policy is used only if you enable IKEv2 in the VPN in the IKE Proposal and IPsec Proposal policies, and if at least some of the devices in the topology support IKEv2.

To configure IKEv2, the device must be an ASA running ASA Software release 8.4(1) or higher. For more information on device support, see Understanding Devices Supported by Each IPsec Technology, page 25-9.

Tip If you support only IKEv2 in the topology, ensure that you unassign the IKEv1 Preshared Keys and IKEv1 Public Key Infrastructure policies to avoid validation warnings.

Related Topics

- Understanding IKE, page 26-5
- Deciding Which Authentication Method to Use, page 26-8

Step 1 Open the Site-to-Site VPN Manager Window, page 25-18, select a regular IPsec topology (that supports IKEv2) in the VPNs selector, then select **IKEv2 Authentication** in the Policies selector.

For reference information on the policy, see IKEv2 Authentication Policy, page 26-70.

Step 2 On the **Global IKEv2 Authentication Settings** tab, configure the authentication type that should be used for devices in the VPN for which no override is configured on the Override IKEv2 Authentication Settings tab. Select the option that is used by most devices in the VPN. You can configure a global preshared key or trustpoint:

- **Global Preshared Keys**—To configure a global preshared key, select **Key Specification** and then configure one of the following options:
  - User Defined—Enter the desired global key and enter it again in the Confirm field.
  - Auto Generated—Enter the length of the key that should be generated and select whether you want to use the same key for all tunnels or the same key at both ends of a single tunnel. If you select neither of these options, unique keys are generated for every end point.

    You can also select **Regenerate Key (On Next Deployment)** to have new keys generated. This allows you to periodically re-key the VPN. The check box is cleared after the next successful deployment.
- **Global Trustpoint (CA Servers)**—To configure trustpoint certificate authorization, select **PKI Specification** and enter the name of the PKI enrollment object that identifies the Certificate Authority (CA) server.

  Note Ensure that you enter the same object name as deployed in the PKI policy (see Step 2 of Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56).

  Click **Select** to select the object from a list or to create a new object.

Step 3 If you want to override the global IKEv2 authentication configuration for specific devices, click the **Override IKEv2 Authentication Settings** tab and do any of the following:

- To add an override, click the **Add Row (+)** button and fill in the IKEv2 Authentication dialog box. You select the local and remote peers for which to create the override, and then specify the preshared key or CA server that should be used. See IKEv2 Authentication (Override) Dialog Box, page 26-72.
- To edit an override, select it in the table and click the **Edit Row (pencil)** button.
- To delete an override, select it in the table and click the **Delete Row (trash can)** button.

Note You can configure asymmetric authentication in site-to-site VPNs, where each side of a tunnel can have different preshared keys. To create asymmetric keys for IKEv2 authentication, for each peer device that is part of the site-to-site topology, you must add two rows on the Override IKEv2 Authentication Settings tab. For more information, see IKEv2 Authentication (Override) Dialog Box, page 26-72.

IKEv2 Authentication Policy

Use the IKEv2 Authentication policy to configure the device authentication settings for Internet Key Exchange (IKE) version 2 in site-to-site VPNs. These settings apply to ASA 8.4(1)+ devices only. For more information about configuring IKEv2 authentication, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

The policy contains two tabs:

- Global IKEv2 Authentication Settings—The global settings apply to all devices in the VPN unless overrides are configured on the Overrides tab. Configure the global settings to represent the authentication scheme used by most devices in the VPN.
- Override IKEv2 Authentication Settings—The override settings apply unique authentication settings to specific tunnels, allowing you to create unique preshared key and trustpoint combinations that are required by various tunnels in the VPN. The settings you configure on this tab are used first and always take precedence over the global settings.

Navigation Path

Open the Site-to-Site VPN Manager Window, page 25-18, select a regular IPsec topology (that supports IKEv2) in the VPNs selector, then select **IKEv2 Authentication** in the Policies selector.

This policy is not available as a shared policy.

Related Topics

- Understanding IKE, page 26-5
- Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19
- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49

Field Reference

Table 26-15 IKEv2 Authentication Policy

Element	Description
Global IKEv2 Authentication Settings Tab	
Key Specification	Use a preshared key for authentication in the VPN. Configure one of the following: • **User Defined**—Enter the desired global key and enter it again in the Confirm field. The key can be 1 to 128 characters. • **Auto Generated**—Have Security Manager generate a key for you. Specify the following options to indicate how the key should be generated: – **Key Length**—The length of the key that should be generated, from 1 to 128. – **Same Keys for All Tunnels**—Select this option to generate the same keys for all tunnels in the VPN. If you do not select this option, different keys or pair of keys (if you select Same Key for Tunnel Endpoints) are used for each tunnel. – **Same Key for Tunnel Endpoints**—Select this option to generate the same key on each end of each tunnel within the VPN. If you do not select this option, different keys are generated on each end of the tunnel. – **Regenerate Key (On Next Deployment)**—Select this option to generate new keys for the next deployment to the devices. This allows you to easily re-key the VPN. After a successful deployment, this check box is cleared so that keys are not regenerated on the subsequent deployment. Select the option each time you want to re-key the VPN.
PKI Specification	The name of the PKI enrollment policy object that defines the trustpoint for IKEv2 connections. A trustpoint represents a Certificate Authority (CA)/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate. Click **Select** to select the PKI enrollment object or to create a new object.

Table 26-15 IKEv2 Authentication Policy (continued)

Element	Description
Override IKEv2 Authentication Settings tab	The table lists the IKEv2 authentication overrides defined for the VPN. These policies take precedence over the preshared key/PKI configuration defined in the global settings. Do any of the following to configure overrides: • To add an override, click the **Add Row (+)** button and fill in the IKEv2 Authentication dialog box. You select the local and remote peers for which to create the override, and then specify the preshared key or CA server that should be used. See IKEv2 Authentication (Override) Dialog Box, page 26-72. • To edit an override, select it in the table and click the **Edit Row (pencil)** button. • To delete an override, select it in the table and click the **Delete Row (trash can)** button. **Note** You can configure asymmetric authentication in site-to-site VPNs, where each side of a tunnel can have different preshared keys. To create asymmetric keys for IKEv2 authentication, for each peer device that is part of the site-to-site topology, you must add two rows on the Override IKEv2 Authentication Settings tab. For more information, see IKEv2 Authentication (Override) Dialog Box, page 26-72

IKEv2 Authentication (Override) Dialog Box

Use the IKEv2 Authentication dialog box to configure overrides to the IKEv2 authentication global settings for a site-to-site VPN. For more information about IKEv2 global and override authentication settings, see Configuring IKEv2 Authentication in Site-to-Site VPNs, page 26-68.

Navigation Path

From the Override IKEv2 Authentication Settings tab of the IKEv2 Authentication policy (see IKEv2 Authentication Policy, page 26-70), click the **Add Row (+)** button or select an override in the table and click **Edit Row (pencil)**.

Field Reference

Table 26-16 IKEv2 Authentication Dialog Box

Element	Description
Local Peers Remote Peers	The local and remote sides of the tunnels for which you are defining this override. To add devices to the list, click the **Select** button to the right of the list to open the Local or Remote Peer Selection dialog box. In that dialog box, select the desired peers in the Available list and click **>>** to move them to the Selected list. You can deselect a device by doing the reverse (using the **<<** button). The list of available devices includes only those devices that support IKEv2 connections, which might not be all of the devices in the VPN.
IKEv2 Authentication Mode	The IKEv2 authentication mode to use between the selected local and remote peers. Select one of the following: • Key Specification—A user-defined preshared key, from 1 to 128 characters. Enter the desired key and enter it again in the Confirm field. • PKI Specification—The name of the PKI enrollment policy object that defines the trustpoint for IKEv2 connections. Click **Select** to select the PKI enrollment object or to create a new object.

Configuring Asymmetric keys for IKEv2 Authentication

You can configure asymmetric authentication in site-to-site VPNs, where each side of a tunnel can have different preshared keys. To create asymmetric keys for IKEv2 authentication, for each peer device that is part of the site-to-site topology, you must add two rows on the Override IKEv2 Authentication Settings tab. Perform the following:

1. Click the **Override IKEv2 Authentication Settings** tab and then click the**Add Row (+)** button. The IKEv2 Authentication dialog box opens. On Peers specification, select the Local Peer device and the Remote Peer device that are part of the site-to-site VPN topology. On **IKEv2 Authentication Mode** select Key Specification and then specify a key and confirm. Security Manager considers this key as the local preshared key for the selected local peer device and the same key as the remote preshared key for the selected remote peer device. Click **OK** to return to the Override IKEv2 Authentication Settings tab.
2. With the **Override IKEv2 Authentication Settings** tab selected, click the **Add Row (+)** button. The IKEv2 Authentication dialog box opens. On Peers specification, for Local Peer, select the Remote Peer device of Step 1 and for Remote Peer, select the Local Peer device of Step 1. On **IKEv2 Authentication Mode** select Key Specification and then specify a key and confirm. This key must be different from the key that you specified in Step 1.

   The following table illustrates the configuration of asymmetric keys for IKEv2 authentication:

	Local Peer Device	Remote Peer Device	Authentication method (Preshared Key)
Add Row 1	Peer1	Peer2	test123
Add Row 2	Peer2	Peer1	sample123

GRE and DM VPNs

You can configure Generic Routing Encapsulation (GRE) and Dynamic Multipoint (DM) VPNs that include GRE mode configurations. You can configure IPsec GRE VPNs for hub-and-spoke, point-to-point, and full mesh VPN topologies. DMVPN is available for hub-and-spoke topologies only.

This chapter contains the following topics:

- Understanding the GRE Modes Page, page 27-1
- GRE and Dynamic GRE VPNs, page 27-2
- Dynamic Multipoint VPNs (DMVPN), page 27-9

Understanding the GRE Modes Page

Use the GRE Modes page to define the routing and tunnel parameters for IPsec tunneling with GRE, GRE Dynamic IP, and DMVPN policies.

The content of the policy differs depending on how you access it:

- (Site-to-Site VPN Manager Window, page 25-18) When you select a GRE VPN or DMVPN, the GRE Modes policy contains the properties related to the technology and technology type used in the VPN.
- (Policy view) When you select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy, there is an additional field in the policy called **GRE Method**. From the GRE Method list, you must select the VPN technology and technology type for which you are defining the policy: IPsec/GRE, GRE Dynamic IP, DMVPN, or Large Scale DMVPN. This option controls which fields are displayed in the policy. You cannot change the GRE Method after you save the policy.

  When you assign a shared GRE Modes policy to a VPN, the GRE Method and the VPN's technology and type must match or the policy cannot be selected. For example, you cannot assign a shared DMVPN GRE Modes policy to an IPsec/GRE VPN.

The following topics describe the GRE Modes policy in detail based on the selected GRE Methods:

- IPsec/GRE or GRE Dynamic IP—See Configuring GRE Modes for GRE or GRE Dynamic IP VPNs, page 27-6.
- DMVPN or Large Scale DMVPN—See Configuring GRE Modes for DMVPN, page 27-12.

Note When configuring an IPsec/GRE, GRE Dynamic IP, or DMVPN routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Related Topics

- Understanding GRE, page 27-2
- Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5
- Understanding DMVPN, page 27-10
- Understanding IPsec Technologies and Policies, page 25-5

GRE and Dynamic GRE VPNs

You can use Generic Routing Encapsulation (GRE) to create VPNs using Cisco IOS security routers and Catalyst 6500/7600 devices in hub-and-spoke, point-to-point, and full mesh VPN topologies.

This section contains the following topics:

- Understanding GRE, page 27-2
- Configuring IPsec GRE VPNs, page 27-5
- Configuring GRE Modes for GRE or GRE Dynamic IP VPNs, page 27-6

Understanding GRE

Generic Routing Encapsulation (GRE) is a tunneling protocol that encapsulates a variety of protocol packet types inside IP tunnels, creating a virtual point-to-point connection to devices at remote points over an IP network. With this technology, GRE encapsulates the entire original packet with a standard IP header and GRE header before the IPsec process. Then, IPsec views the GRE packet as an unremarkable IP packet and performs encryption and authentication services, as dictated by the IKE negotiated parameters. Because GRE can carry multicast and broadcast traffic, it is possible to configure a routing protocol for virtual GRE tunnels. The routing protocol detects loss of connectivity and reroutes packets to the backup GRE tunnel, thus providing high resiliency.

For VPN resilience, a spoke must be configured with two GRE tunnels, one to the primary hub and the other to the backup hub. Both GRE tunnels are secured with IPsec: each one has its own IKE security association (SA) and a pair of IPsec SAs. An associated routing protocol automates the failover mechanism, transferring to the backup tunnel if virtual link loss is detected.

Note GRE can be configured on Cisco IOS security routers and Catalyst 6500/7600 devices in hub-and-spoke, point-to-point, and full mesh VPN topologies.

This section contains the following topics:

- Advantages of IPsec Tunneling with GRE, page 27-3
- How Does Security Manager Implement GRE?, page 27-3
- Prerequisites for Successful Configuration of GRE, page 27-3

- Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5

Advantages of IPsec Tunneling with GRE

The main advantages of IPsec tunneling with GRE are the following:

- GRE uses a routing protocol by which every IPsec peer knows the status of every other peer at all times.
- GRE provides higher resiliency than IKE keepalive.
- Spoke-to-spoke connectivity is supported when you use GRE.
- GRE supports multicast and broadcast transmissions.

Note GRE does not support the use of dynamic cryptographic tunnels.

How Does Security Manager Implement GRE?

Security Manager implements an additional Interior Gateway Protocol (IGP) solution for GRE. An IGP refers to a group of devices that receive routing updates from one another by a routing protocol, EIGRP, OSPF, or RIP. Each "routing group" is identified by a logical number. For general routing purposes, the interfaces on the routers in your networks belong to an IGP. Security Manager adds an additional IGP that is dedicated for IPsec and GRE-secured communication. This additional IGP is the secured IGP. The existing IGP (unsecured IGP), is used for routing traffic that does not require encryption.

For a GRE tunnel to be established, Security Manager configures a virtual interface on each device. These virtual interfaces are the endpoints of the GRE tunnel. Each virtual interface is unique. The GRE tunnel interface has an IP address (inside tunnel IP address) which is taken from an interface that Security Manager creates. The GRE tunnel points to the source and destination IP addresses of either the physical or loopback interfaces on each device. The GRE virtual interfaces belong to the secured IGP, as do the inside interfaces. Routing updates within the secured IGP are GRE encapsulated and IPsec is applied. A flow whose destination is a secured interface (according to the routing updates of the secured IGP) is directed through the GRE interface where it is GRE encapsulated and then evaluated against the crypto ACL. If it matches the crypto ACL, it is routed through the GRE and VPN tunnels.

Prerequisites for Successful Configuration of GRE

Consider the following prerequisites before using GRE in your network:

- You must identify the inside interfaces on your devices—the physical interfaces on the device that connect the device with its internal subnets and networks.
- You must select a routing protocol (known as an IGP) or a static route, whenever you enable GRE.

  Security Manager supports the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes.

  - EIGRP—Enhanced Interior Gateway Routing Protocol enables the exchange of routing information within an autonomous system and addresses some of the more difficult issues associated with routing in large, heterogeneous networks. Compared to other protocols, EIGRP provides superior convergence properties and operating efficiency, and combines the advantages of several different protocols. For more information, see EIGRP Routing on Cisco IOS Routers, page 67-8.

  - OSPF—Open Shortest Path First is a link-state, hierarchical protocol that features least-cost routing, multipath routing, and load balancing.

    Using OSPF, a host that obtains a change to a routing table or detects a change in the network immediately multicasts the information to all other hosts in the network, so that all will have the same routing table information. For more information, see OSPF Routing on Cisco IOS Routers, page 67-19.
  - RIPv2—Routing Information Protocol is a distance-vector protocol that sends routing-update messages at regular intervals and whenever the network topology changes.

    Using RIPv2, a gateway host (with a router) sends its entire routing table to its closest neighbor host every 30 seconds, which in turn passes the information on to its next neighbor, and so on, until all hosts within the network have the same knowledge of routing paths. RIPv2 uses a hop count to determine network distance. Each host with a router in the network uses the routing table information to determine the next host to route a packet to for a specified destination.

    RIP is considered an effective solution for small homogeneous networks. For larger, more complicated networks, RIP's transmission of the entire routing table every 30 seconds may put a heavy amount of extra traffic in the network. For more information, see RIP Routing on Cisco IOS Routers, page 67-42.
  - Static route—Use a static routing policy to provide a robust, stable IPsec-protected GRE tunnel if there is a fixed, unchanging route between two devices. For each device subnet, a static route is created on the device pointing to the corresponding tunnel interface. For more information, see Static Routing on Cisco IOS Routers, page 67-50.

- You must specify an IGP process number. The IGP process number identifies the IGP process to which the inside interface on the device belongs. When GRE is implemented, this will be the secured IGP. For secure communication, the inside interfaces on the devices in your VPN must use the same IGP process. The IGP process number must be within a specified range. If you have an existing IGP process on the device that is within this range, but is different from the IGP process number specified in your GRE settings, Security Manager removes the existing IGP process. If the existing IGP process matches the one specified in your GRE settings, any networks included in the existing IGP process that do not match the specified inside interfaces are removed.
- If the inside interfaces on your devices are configured to use an IGP process other than the IGP process specified in your GRE settings (meaning that the interfaces belong to an unsecured IGP):
  - For spokes: Manually remove the inside interfaces from the unsecured IGP through the device CLI before configuring GRE.
  - For hubs: If the hub inside interface is used as a network access point for Security Manager, then on deployment, the interface is advertised in both secured and unsecured IGPs. To ensure that the spoke peers use only the secured IGP, manually add the auto-summary command for the unsecured IGP or remove the unsecured IGP for that inside interface.
- You must provide a subnet that is unique yet it can be non-globally-routable for loopback. This subnet must only be used to support the implementation of loopback for GRE. The loopback interfaces are created, maintained, and used only by Security Manager. You should not use them for any other purpose.
- If you are using static routes, not unsecured IGP, make sure you configure static routes on the spokes through to the hub inside interfaces.

Note You can configure the above settings in the GRE Modes page when IPsec/GRE is the selected IPsec technology.

Understanding GRE Configuration for Dynamically Addressed Spokes

When a spoke has a dynamic IP address, there is no fixed GRE tunnel source address (to be used by the GRE tunnel on the spoke side) or destination address (to be used by the GRE tunnel on the hub side). Therefore, Security Manager creates additional loopback interfaces on the hub and the spoke, to be used as the GRE tunnel endpoints. You must specify a subnet from which Security Manager can allocate an IP address for the loopback interfaces.

Note GRE Dynamic IP can only be configured on Cisco IOS routers and Catalyst 6500/7600 devices in hub-and-spoke VPN topologies.

Security Manager uses the Cisco Configuration Engine to retrieve device IP addresses and other information from dynamically addressed devices. Devices that have dynamic IP addresses connect to the Configuration Engine manager at periodic intervals to upgrade device configuration files and to pass device and status information.

For more information, see Adding, Editing, or Deleting Auto Update Servers or Configuration Engines, page 3-36.

Note You can configure the GRE Dynamic IP settings in the GRE Modes page when GRE Dynamic IP is the selected IPsec technology.

Related Topics

- Understanding GRE, page 27-2
- Configuring GRE Modes for GRE or GRE Dynamic IP VPNs, page 27-6

Configuring IPsec GRE VPNs

To configure an IPsec GRE (generic routing encapsulation) VPN, use the Create VPN wizard as described in Creating or Editing VPN Topologies, page 25-28. You can also edit the membership of the VPN, or some of its policies, using the described procedures. If you are creating a hub-and-spoke VPN with dynamically addressed spokes, also see Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5.

If you need to make changes to other policies and settings, open the policies from the Site-to-Site Manager page, as follows:

- For ISAKMP and IPSec settings, select **VPN Global Settings**. See Configuring VPN Global Settings, page 26-30.
- For IKE proposal policies, select **IKE Proposal**. See Configuring an IKE Proposal, page 26-9.
- For IPSec proposals, select **IPsec Proposal**. See Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22.
- For preshared key policies, select **IKEv1 Preshared Key**. See Configuring IKEv1 Preshared Key Policies, page 26-48.
- For public key (PKI) policies, select **Public Key Infrastructure**. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.
- For Generic Routing Encapsulation configuration, select **GRE Modes**. See Configuring GRE Modes for DMVPN, page 27-12.

Related Topics

- Understanding IKE, page 26-5
- Understanding GRE, page 27-2
- Prerequisites for Successful Configuration of GRE, page 27-3
- Advantages of IPsec Tunneling with GRE, page 27-3

Configuring GRE Modes for GRE or GRE Dynamic IP VPNs

Use the GRE Modes policy to define the routing and tunnel parameters for IPsec tunneling in a GRE or GRE Dynamic IP VPN.

To open the GRE Modes policy:

- (Site-to-Site VPN Manager Window, page 25-18) Select an IPsec/GRE or GRE Dynamic IP topology, then select **GRE Modes** from the policies list.
- (Policy view) Select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy. Then, select either IPsec/GRE or Dynamic GRE from the **GRE Method** list.

The following table describes the elements on the GRE Modes page for configuring IPsec tunneling with GRE or GRE Dynamic IP.

Note When configuring a GRE routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Table 27-1 GRE Modes Page for GRE or GRE Dynamic IP VPNs

Element	Description
Routing Parameters Tab	
Routing Protocol	Select the required dynamic routing protocol (EIGRP, OSPF, or RIPv2,) or static route to be used for GRE or GRE Dynamic IP. The default routing protocol is EIGRP. For more information about configuring these protocols, see Prerequisites for Successful Configuration of GRE, page 27-3.
AS Number (EIGRP only.)	The number that is used to identify the autonomous system (AS) area to which the EIGRP packet belongs. The range is 1-65535. The default is 110. An autonomous system (AS) is a collection of networks that share a common routing strategy. An AS can be divided into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. An AS ID identifies the area to which the packet belongs. All EIGRP packets are associated with a single area, so all devices must have the same AS number.
Hello Interval (EIGRP only.)	The interval between hello packets sent on the interface, between 1 and 65535 seconds. The default is 5 seconds.

Table 27-1 GRE Modes Page for GRE or GRE Dynamic IP VPNs (continued)

Element	Description
Hold Time (EIGRP only.)	The number of seconds the router will wait to receive a hello message before invalidating the connection. The range is between 1 and 65535. The default hold time is 15 seconds (three times the hello interval).
Delay (EIGRP only.)	The throughput delay for the primary route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1000.
Failover Delay (EIGRP only.)	The throughput delay for the failover route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1500.
Bandwidth (EIGRP only.)	The amount of bandwidth available to the primary route interface for the EIGRP packets. You should enter a value that gives priority to the primary route over other routes. You can enter a value in the range 1 to 10000000 kb. The default is 1000 kb. **Note** By default, the cost of sending a packet on an interface is calculated based on the bandwidth—the higher the bandwidth, the lower the cost.
Failover Bandwidth (EIGRP only.)	The amount of bandwidth available to the failover route interface for the EIGRP packets. Enter a value in the range 1 to 10000000 kb. The default is 1000 kb.
Process Number (OSPF only.)	The routing process ID number that will be used to identify the secured IGP that Security Manager adds when configuring GRE. The range is between 1 and 65535. The default is 110. Security Manager adds an additional Interior Gateway Protocol (IGP) that is dedicated for IPsec and GRE secured communication. An IGP refers to a group of devices that receive routing updates from one another by means of a routing protocol. Each “routing group” is identified by the process number. For more information, see Understanding GRE, page 27-2.
Hub Network Area ID (OSPF only.)	The ID number of the area in which the hub’s protected networks will be advertised, including the tunnel subnet. You can specify any number. The default is 0.
Spoke Protected Network Area ID (OSPF only.)	The ID number of the area in which the remote protected networks will be advertised, including the tunnel subnet. You can specify any number. The default is 1.
Authentication (OSPF or RIPv2 only.)	A string that specifies the OSPF or RIPv2 authentication key. The string can be up to eight characters long.
Cost (OSPF or RIPv2 only.)	The cost of sending a packet on the primary route interface. If the selected protocol is OSPF, enter a value in the range 1-65535; the default is 100. If the selected protocol is RIPv2, enter a value in the range 1-15; the default is 1.

Table 27-1 GRE Modes Page for GRE or GRE Dynamic IP VPNs (continued)

Element	Description
Failover Cost (OSPF or RIPv2 only.)	The cost of sending a packet on the secondary (failover) route interface. You can enter a value in the range 1-65535 for OSPF (the default is 125), or in the range 1-15 for RIPv2 (the default is 2).
Filter Dynamic Updates on Spokes	When selected, enables the creation of a redistribution list that filters all dynamic routing updates on the spokes. This forces the spoke devices to advertise (populate on the hub device) only their own protected subnets and not other IP addresses.
Tunnel Parameters Tab	
Tunnel IP	Select the required option to specify the GRE or GRE Dynamic IP tunnel interface IP address. • **Use Physical Interface**—When selected, uses the private IP address of the tunnel taken from the protected network. • **Use Subnet**—When selected, uses the tunnel IP address taken from an IP range. This is the default. In the Subnet field, enter the private IP address including the unique subnet mask (default is 1.1.1.0/24). If you are also configuring a dial backup interface, enter its subnet in the Dial Backup Subnet field provided (default is 1.1.2.0/24). **Note** In most cases, when you use a subnet to specify a GRE tunnel interface IP address, Security Manager creates a loopback interface on the device which is used for the tunnel IP address. If the device belongs to a VPN topology whose configurations were discovered by Security Manager, and you configure an IP address directly on the device's GRE tunnel, Security Manager keeps that configuration and does not create a loopback interface on the device. However, a loopback is always configured on a hub in a VPN topology; in a hub-and-spoke VPN topology with multiple hubs, a loopback interface is also configured on the spokes. • **Use Loopback Interface**—When selected, uses the tunnel IP address taken from an existing loopback interface. In the Role field, enter the name of the interface role object that defines the loopback interface name, or click **Select** to select it from a list or to create a new object. **Note** To view the new GRE tunnel or loopback interfaces in the Router Interfaces page, you must rediscover the device inventory details after successfully deploying the VPN to the device.

Table 27-1 GRE Modes Page for GRE or GRE Dynamic IP VPNs (continued)

Element	Description
Configure Unique Tunnel Source for each Tunnel	When enabled, each GRE tunnel interface in the VPN is assigned a unique tunnel source. In the **Tunnel Source IP Range** field, enter a subnet IP to be used as tunnel sources. **Note** When enabled, this feature is set for all GRE tunnel interfaces in the VPN. If you want to assign a specific tunnel source for an interface, use the Peers policy to configure the endpoints for the desired devices; see Defining the Endpoints and Protected Networks, page 25-34.
Tunnel Source IP Range (GRE Dynamic IP only.)	The private IP address including the unique subnet mask that supports the loopback for GRE. The GRE tunnel interface has an IP address (inside tunnel IP address) which is taken from a loopback interface that Security Manager creates specifically for this purpose. When a spoke has a dynamic IP address, there is no fixed GRE tunnel source address (to be used by the GRE tunnel on the spoke side) or destination address (to be used by the GRE tunnel on the hub side). Therefore, Security Manager creates additional loopback interfaces on the hub and the spoke to use as the GRE tunnel endpoints. You must specify a subnet from which Security Manager can allocate an IP address for the loopback interfaces.
Enable IP Multicast	When selected, enables multicast transmissions across your GRE tunnels. IP multicast delivers application source traffic to multiple receivers without burdening the source or the receivers, while using a minimum of network bandwidth.
Rendezvous Point	Only available if you selected the Enable IP Multicast check box. If required, you can enter the IP address of the interface that will serve as the rendezvous point (RP) for multicast transmission. Sources send their traffic to the RP. This traffic is then forwarded to receivers down a shared distribution tree.

Dynamic Multipoint VPNs (DMVPN)

Dynamic Multipoint VPN (DMVPN) is a hub-and-spoke VPN technology that enables better scaling of large and small IPsec VPNs by combining generic routing encapsulation (GRE) tunnels, IP Security (IPsec) encryption, and Next Hop Resolution Protocol (NHRP) routing.

This section contains the following topics:

- Understanding DMVPN, page 27-10
- Configuring DMVPN, page 27-12
- Configuring GRE Modes for DMVPN, page 27-12
- Configuring Large Scale DMVPNs, page 27-16
- Configuring Server Load Balancing in Large Scale DMVPN, page 27-17

Understanding DMVPN

Dynamic Multipoint VPN (DMVPN) enables better scaling of large and small IPsec VPNs by combining generic routing encapsulation (GRE) tunnels, IP Security (IPsec) encryption, and Next Hop Resolution Protocol (NHRP) routing. (For information about large scale DMVPNs, see Configuring Large Scale DMVPNs, page 27-16.)

Security Manager supports DMVPN using the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes. In addition, On-Demand Routing (ODR) is supported. ODR is not a routing protocol. It may be used in a hub-and-spoke VPN topology when the spoke routers do not connect to any router other than the hub. If you are running dynamic protocols, ODR is not suitable for your network environment.

You can use DMVPN on a hub-and-spoke VPN topology only with devices running Cisco IOS Software release 12.3T devices and later, or ASRs running Cisco IOS XE Software 2.x or later (known as 12.2(33)XNA+ in Security Manager). DMVPN is not supported on Catalyst VPN Services Module devices or on High Availability (HA) groups. If your device does not support DMVPN, use GRE dynamic IP to configure GRE for dynamically addressed spokes. See Understanding GRE Configuration for Dynamically Addressed Spokes, page 27-5.

The following topics provide more overview information on DMVPN:

- Enabling Spoke-to-Spoke Connections in DMVPN Topologies, page 27-10
- Advantages of DMVPN with GRE, page 27-11

The following documents on Cisco.com explain DMVPN in further detail:

- *Cisco Dynamic Multipoint VPN: Simple and Secure Branch-to-Branch Communications*—Explains DMVPN technology and where and why you would use it. This data sheet explains the technologies used with DMVPN and the benefits derived from those technologies.
- *Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3*—Explains the difference between phase 2 and phase 3 spoke-to-spoke connections. Creating spoke-to-spoke connections is a configuration option with DMVPN. Phase 3 uses shortcut switching enhancements to increase network performance and scalability.
- Additional white papers and presentations are available at http://www.cisco.com/en/US/products/ps6658/prod_literature.html.

Enabling Spoke-to-Spoke Connections in DMVPN Topologies

You can use DMVPN to essentially create a full-mesh VPN, in which traditional hub-and-spoke connectivity is supplemented by dynamically-created IPsec tunnels directly between the spokes. With direct spoke-to-spoke tunnels, traffic between remote sites does not need to traverse the hub; this eliminates additional delays and conserves WAN bandwidth. Spoke-to-spoke capability is supported in a single-hub or multihub environment. Multihub deployments provide increased spoke-to-spoke resiliency and redundancy.

You can use the 80:20 traffic rule to determine whether to use a pure hub-and-spoke topology or to allow direct spoke-to-spoke connections:

- If 80 percent or more of the traffic from the spokes are directed into the hub network itself, deploy the hub-and-spoke model.
- If more than 20 percent of the traffic is meant for other spokes, consider the spoke-to-spoke model.

For networks with a high volume of IP Multicast traffic, the hub-and-spoke model is usually preferred.

When you configure the GRE Modes policy for a DMVPN, you can elect to allow spokes to create these direct connections. You must select the DMVPN phase to use for these connections:

- **Phase 2**—Spoke to spoke connections go through regional hubs and routing protocol updates from hubs to spokes are not summarized.
- **Phase 3** (Default)—Spokes can create direct connections with each other and routing updates from hubs to spokes are summarized. This option allows the greatest scalability and reduces latency. Devices must run IOS Software release 12.4(6)T or higher; ASRs must run IOS XE Software release 2.4 (called 12.2(33)XND) or higher. Security Manager automatically creates a phase 2 configuration for devices running a lower OS version.

For more information on configuring the GRE Modes policy, see Configuring GRE Modes for DMVPN, page 27-12.

Related Topics

- Understanding DMVPN, page 27-10
- *Cisco Dynamic Multipoint VPN: Simple and Secure Branch-to-Branch Communications*
- *Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3*

Advantages of DMVPN with GRE

Using DMVPN with GRE provides the following advantages:

- **Simplified GRE configuration on the hub**

  With GRE, a tunnel is configured on the hub for each connected spoke. With GRE + DMVPN, only one tunnel is configured for all the connected spokes.

- **Support for dynamically addressed spokes**

  When using GRE, the physical interface IP address of the spoke routers must be configured as the GRE tunnel destination address, when configuring the hub router. DMVPN enables spoke routers to have dynamic external interface IP addresses, and provides robust configuration that does not have to be redeployed to the device even if the external interface IP address changes. When the spoke comes online, it sends to the hub registration packets that contain the physical interface IP address of the spoke.

- **Dynamic tunnel creation for direct spoke-to-spoke communication**

  NHRP enables spoke routers to dynamically learn the external interface IP address of the routers in the VPN network. Using NHRP, the hub maintains an NHRP database of the public interface addresses of all the spokes (the clients). Each spoke registers its real address with the hub when it boots.

  When a spoke wants to transmit a packet to another spoke, it can use NHRP to dynamically determine the required destination address of the destination spoke. The hub acts as the NHRP server, handling the request for the source spoke. This enables the dynamic creation of an IPsec+GRE tunnel directly between spoke routers, without having to go through a hub router, thus reducing the delay of multiple encryption and decryption actions on the hub.

Configuring DMVPN

To configure a hub-and-spoke Dynamic Multipoint VPN, use the Create VPN wizard as described in Creating or Editing VPN Topologies, page 25-28. You can also edit the membership of the VPN, or some of its policies, using the described procedures. If you are creating a Large Scale DMVPN, also see Configuring Large Scale DMVPNs, page 27-16.

If you need to make changes to other policies and settings, open the policies from the Site-to-Site Manager page, as follows:

- For ISAKMP and IPSec settings, select **VPN Global Settings**. See Configuring VPN Global Settings, page 26-30.
- For IKE proposal policies, select **IKE Proposal**. See Configuring an IKE Proposal, page 26-9.
- For IPSec proposals, select **IPsec Proposal**. See Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22.
- For preshared key policies, select **IKEv1 Preshared Key**. See Configuring IKEv1 Preshared Key Policies, page 26-48.
- For public key (PKI) policies, select **Public Key Infrastructure**. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.
- For Generic Routing Encapsulation configuration, including the selection of phase 2 or 3 connections between spokes, select **GRE Modes**. See Configuring GRE Modes for DMVPN, page 27-12.
- For server load balancing policies that are used with Large Scale DMVPN, select **Server Load Balance**. See Configuring Server Load Balancing in Large Scale DMVPN, page 27-17.

Related Topics

- Understanding IKE, page 26-5
- Understanding DMVPN, page 27-10
- Enabling Spoke-to-Spoke Connections in DMVPN Topologies, page 27-10
- Advantages of DMVPN with GRE, page 27-11

Configuring GRE Modes for DMVPN

Use the GRE Modes policy to define the routing and tunnel parameters for IPsec tunneling in a DMVPN.

To open the GRE Modes policy:

- (Site-to-Site VPN Manager Window, page 25-18) Select a DMVPN or Large Scale DMVPN topology, then select **GRE Modes** from the policies list.
- (Policy view) Select **Site-to-Site VPN > GRE Modes**, and create a new policy or select an existing policy. Then, select either DMVPN or Large Scale DMVPN from the **GRE Method** list.

The following table describes the elements on the GRE Modes page for configuring a DMVPN.

Note When configuring a DMVPN routing policy, Security Manager adds a routing protocol to all the devices in the secured IGP, on deployment. If you want to maintain this secured IGP, you must create a router platform policy (on each member device) using the same routing protocol and autonomous system (or process ID) number as defined in the GRE Modes policy.

Table 27-2 GRE Modes Page for DMVPN

Element	Description
Routing Parameters Tab	
Routing Protocol	Select the required dynamic routing protocol, or static route, to be used in the DMVPN tunnel. Options include the EIGRP, OSPF, and RIPv2 dynamic routing protocols, and GRE static routes. On-Demand Routing (ODR) is also supported. On-Demand Routing is not a routing protocol. It can be used in a hub-and-spoke VPN topology when the spoke routers connect to no other router other than the hub. If you are running dynamic protocols, On-Demand Routing is not suitable for your network environment. For more information, see Understanding GRE, page 27-2.
AS Number (EIGRP only.)	The number that is used to identify the autonomous system (AS) area to which the EIGRP packet belongs. The range is 1-65535. The default is 110. An autonomous system (AS) is a collection of networks that share a common routing strategy. An AS can be divided into a number of areas, which are groups of contiguous networks and attached hosts. Routers with multiple interfaces can participate in multiple areas. An AS ID identifies the area to which the packet belongs. All EIGRP packets are associated with a single area, so all devices must have the same AS number.
Hello Interval (EIGRP only.)	The interval between hello packets sent on the interface, from 1 to 65535 seconds. The default is 5 seconds.
Hold Time (EIGRP only.)	The number of seconds the router will wait to receive a hello message before invalidating the connection. The range is 1-65535. The default hold time is 15 seconds (three times the hello interval)
Delay (EIGRP only.)	The throughput delay for the primary route interface, in microseconds. The range of the tunnel delay time is 1-16777215. The default is 1000.
Bandwidth (EIGRP only.)	The bandwidth for the primary route interface, in kilobits. The range of bandwidth is 1 to 10000000. The default is 1000.
Bandwidth (EIGRP only.)	The amount of bandwidth available to the primary route interface for the EIGRP packets. You should enter a value that gives priority to the primary route over other routes. You can enter a value in the range 1 to 10000000 kb. The default is 1000 kb. **Note** By default, the cost of sending a packet on an interface is calculated based on the bandwidth—the higher the bandwidth, the lower the cost.
Process Number (OSPF only.)	The routing process ID number that will be used to identify the secured IGP that Security Manager adds when configuring DMVPN. The valid range for either protocol is 1-65535. The default is 110.

Table 27-2 GRE Modes Page for DMVPN (continued)

Element	Description
Hub Network Area ID (OSPF only.)	The ID number of the area in which the hub's protected networks will be advertised, including the tunnel subnet. You can enter any number. The default is 0.
Spoke Protected Network Area ID (OSPF only.)	The ID number of the area in which the remote protected networks will be advertised, including the tunnel subnet. You can enter any number. The default is 1.
Authentication Key (OSPF and RIPv2.)	A string that indicates the OSPF or RIPv2 authentication key. The string can be up to eight characters long.
Cost (OSPF and RIPv2.)	The cost of sending a packet on the primary route interface. If the selected protocol is OSPF, enter a value in the range 1-65535; the default is 100. If the selected protocol is RIPv2, enter a value in the range 1-15; the default is 1.
Allow Direct Spoke to Spoke Connectivity	Whether to enable direct communication between spokes without going through the hub. Select the DMVPN phase you want to use, which determines the types of connections that spokes can make: • **Phase 2**—Spoke to spoke connections go through regional hubs and routing protocol updates from hubs to spokes are not summarized. • **Phase 3** (Default)—Spokes can create direct connections with each other and routing updates from hubs to spokes are summarized. This option allows the greatest scalability and reduces latency. Devices must run IOS Software release 12.4(6)T or higher; ASRs must run IOS XE Software release 2.4 (called 12.2(33)XND) or higher. Security Manager automatically creates a phase 2 configuration for devices running a lower OS version. For detailed information on how phase 2 and 3 differ, see "Migrating from Dynamic Multipoint VPN Phase 2 to Phase 3" on Cisco.com. **Note** With direct spoke-to-spoke communication, you must use the Main Mode Address option for preshared key negotiation. For more information, see Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47.
Filter Dynamic Updates On Spokes	Unavailable if you are using On-Demand Routing or a static route for your DMVPN tunnel. When selected, enables the creation of a redistribution list that filters all dynamic routing updates (EIGRP, OSPF, and RIPv2) on spokes. This forces the spoke devices to advertise (populate on the hub device) only their own protected subnets and not other IP addresses.

Table 27-2 GRE Modes Page for DMVPN (continued)

Element	Description
Tunnel Parameters Tab	
Tunnel IP Range	The IP address range of the inside tunnel interface IP address, including the unique subnet mask. This field defines a subnet, such as 10.1.1.0/24. **Note** If Security Manager detects that a tunnel interface IP address already exists on the device, and its IP address matches the tunnel's IP subnet field, it will use that interface as the GRE tunnel.
Dial Backup Tunnel IP Range	If you are configuring a dial backup interface, enter its inside tunnel interface IP address range, including the unique subnet mask. This field defines a subnet.
Server Load Balance	When selected, enables the configuration of load balancing on a Cisco IOS router that serves as a hub in a multiple hubs configuration. Server load balancing optimizes performance in a multiple hubs configuration, by sharing the workload. In this configuration, the DMVPN server hubs share the same tunnel IP and source IP addresses, presenting the appearance of a single device to the spokes in a VPN topology.
Enable IP Multicast	When selected, enables multicast transmissions across your GRE tunnels. IP multicast delivers application source traffic to multiple receivers without burdening the source or the receivers, while using a minimum of network bandwidth.
Rendezvous Point	Only available if you selected the Enable IP Multicast check box. If required, you can enter the IP address of the interface that will serve as the rendezvous point (RP) for multicast transmission. Sources send their traffic to the RP. This traffic is then forwarded to receivers down a shared distribution tree.
Tunnel Key	A number that identifies the tunnel key. The default is 1. The tunnel key differentiates between different multipoint GRE (mGRE) tunnel Non Broadcast Multiple Access (NBMA) networks. All mGRE interfaces in the same NBMA network must use the same tunnel key value. If there are two mGRE interfaces on the same router, they must have different tunnel key values. **Note** To view the newly created tunnel interfaces in the Router Interfaces page for routers that are members of the VPN, you must rediscover the device inventory details after successfully deploying the VPN to the device.
NHRP Parameters	
Network ID	All Next Hop Resolution Protocol (NHRP) stations within one logical Non-Broadcast Multi-Access (NBMA) network must be configured with the same network identifier. Enter a globally unique, 32-bit network identifier within the range of 1 to 4294967295.

Table 27-2 GRE Modes Page for DMVPN (continued)

Element	Description
Hold time	The time, in seconds, that routers will keep information provided in authoritative NHRP responses. The cached IP-to-NBMA address mapping entries are discarded after the hold time expires. The default is 300 seconds.
Authentication	An authentication string that controls whether the source and destination NHRP stations allow intercommunication. All routers within the same network using NHRP must share the same authentication string. The string can be up to eight characters long.

Configuring Large Scale DMVPNs

You can configure DMVPN for large scale deployments that might comprise thousands of spokes. In large scale DMVPN topologies, IPsec Terminators, also referred to as Server Load Balance (SLB) devices, reside between the spokes and the hubs. The hubs must be directly connected to the IPsec Terminator—there can be no other device between them.

The IPsec Terminator, which is a Catalyst 6500/7600 device, performs encryption and decryption while the hubs handle all tasks related to Next Hop Resolution Protocol (NHRP) and multipoint generic routing encapsulation (mGRE). The IPsec Terminator is configured to specifically load balance GRE traffic to the hubs, and is configured with dynamic crypto to accept any spokes with any proxies. When using tunnel protection on spokes, these proxies are automatically set to match GRE traffic. One GRE tunnel is configured on the spokes. All hubs connecting to the same IPsec Terminator will use the same Tunnel IP address, and the tunnel source is the Virtual IP address of the IPsec Terminator.

In Security Manager, you configure a Large Scale DMVPN during the creation of a new hub-and-spoke VPN topology as described in Creating or Editing VPN Topologies, page 25-28. You cannot edit an existing standard DMVPN and convert it to a Large Scale DMVPN. When you create the Large Scale DMVPN, keep the following points in mind:

- When you define the technology of the VPN, select **DMVPN** as the technology, and **Large Scale with IPsec Terminator** as the type. For the procedure, see Defining the Name and IPsec Technology of a VPN Topology, page 25-30.
- When you select the devices for the VPN, select the required IPsec Terminators (Catalyst 6500/7600 devices), the hubs and all the spokes. For the procedure, see Selecting Devices for Your VPN Topology, page 25-32.

  There must be direct connectivity between the IPsec Terminators and the hubs.
- When you configure the endpoints, as described Defining the Endpoints and Protected Networks, page 25-34, configure the following in the Edit Endpoints dialog box:
  - For each hub device, in the Hub Interface tab, select the interface that is connected to the IPsec Terminator. Each hub can be connected to only one IPsec Terminator. Also, identify the protected networks. Each hub in the Large Scale DMVPN must identify itself and its protected networks.
  - For each IPsec Terminator in the Large Scale DMVPN, specify a VPN external interface, the crypto engine slot, and the Inside VLAN. No protected networks are configured on an IPsec Terminator.

After you create the Large Scale DMVPN topology, a Server Load Balance policy is configured on the IPsec Terminators with all the required parameters, which you can edit if required. Initially, all hubs are given the same priority and number of VPN connections. For information on configuring the Server Load Balance policy, see Configuring Server Load Balancing in Large Scale DMVPN, page 27-17.

Note VRF-Aware IPsec cannot be configured in a Large Scale DMVPN.

Related Topics

- Understanding DMVPN, page 27-10
- Configuring DMVPN, page 27-12

Configuring Server Load Balancing in Large Scale DMVPN

Use the Server Load Balance page to view or edit the server load balance policy configured on the IPsec Terminators in a large scale DMVPN. Server load balancing optimizes performance in multiple hub-and-spoke VPN topologies by sharing the workload among a group of hubs. In large scale DMVPN configurations, the IPsec Terminators perform the traffic load balancing. For more information, see Configuring Large Scale DMVPNs, page 27-16.

A weighted round robin (WRR) scheduling algorithm is used to control the bandwidth allocated to output transmission queues. Weighting is based on the amount of bandwidth used by each transmit queue on an interface. Packets from queues with higher capacity are transmitted more often than those from queues with less capacity.

To open the Server Load Balance policy, in the Site-to-Site VPN Manager Window, select an existing Large Scale DMVPN topology, then select **Server Load Balance** from the Policies list.

The table displays the hubs in the VPN, the hub's weight relative to other hubs connected to the same IPsec Terminator, and the maximum number of active connections allowed for the hub. To change the weight or maximum connections, select the hub and click the Edit (pencil) button beneath the table to open the Edit Load Balancing Parameters Dialog Box, page 27-17.

Related Topics

- Configuring Large Scale DMVPNs, page 27-16
- Filtering Tables, page 1-48

Edit Load Balancing Parameters Dialog Box

Use the Edit Load Balancing Parameters dialog box to change the server load balance parameters configured on a hub that is connected to an IPsec Terminator in a large scale DMVPN.

Navigation Path

From the Server Load Balance policy, select a hub and click the **Edit (pencil)** button below the table. For information on opening the Server Load Balance policy, see Configuring Server Load Balancing in Large Scale DMVPN, page 27-17.

Related Topics

- Configuring Large Scale DMVPNs, page 27-16

Field Reference

Table 27-3 Edit Load Balancing Parameters Dialog Box

Element	Description
Weight	The capacity of the hub relative to other hubs connected to the IPsec Terminator, based on the weighted round robin (WRR) scheduling algorithm. You can enter a value between 1 and 255. The default is 1.
Max Connections	The maximum number of active connections to the IPsec Terminator that are permitted to the hub. You can enter a value between 1 and 65535. The default is 500.

Easy VPN

Easy VPN is a hub-and-spoke VPN topology that can be used with a variety of routers, PIX, and ASA devices. Policies are defined mostly on the hub and pushed to remote spoke VPN devices, ensuring that clients have up-to-date policies in place before establishing a secure connection.

This chapter contains the following topics:

- Understanding Easy VPN, page 28-1
- Configuring Client Connection Characteristics for Easy VPN, page 28-7
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Configuring a Connection Profile Policy for Easy VPN, page 28-13
- Configuring a User Group Policy for Easy VPN, page 28-14

Understanding Easy VPN

Easy VPN simplifies VPN deployment for remote offices. With Easy VPN, security policies defined at the head end are pushed to remote VPN devices, ensuring that clients have up-to-date policies in place before establishing a secure connection.

Security Manager supports the configuration of Easy VPN policies on hub-and-spoke VPN topologies. In such a configuration, most VPN parameters are defined on the Easy VPN server, which acts as the hub device. The centrally managed IPsec policies are pushed to the Easy VPN client devices by the server, minimizing the remote (spoke) devices configuration.

The Easy VPN Server can be a Cisco IOS router, a PIX Firewall, or an ASA 5500 series device. The Easy VPN client is supported on PIX 501, 506, 506E Firewalls running PIX 6.3, Cisco 800-3900 Series routers, and ASA 5505 devices running ASA Software release 7.2 or later.

Beginning with version 4.17, Cisco Security Manager provides Easy VPN support with BVI. Typically, Easy VPN determines the highest and lowest security level interfaces during ASA startup. The lowest security level interface is used as the External interface on which vpn client initiates tunnel to the head-end, and highest security level interface is used as Internal Secured interface.

On ASA5506 platform, the default configuration includes BVI with highest security level interface 100 with security level of its member interfaces also set at level 100, along with an external interface with security level 0 (zero). VPN client rejects two or more interfaces having same highest security level. Easy VPN determines that there are more than two interfaces with same highest security level and hence vpn client is not enabled.

In order to overcome this issue, vpnclient secure interface CLI was introduced for all ASA 5506, 5508, and 5512 [x/h/w] devices from ASA 9.9(2) onwards. Thus, to support the CLI in Cisco Security Manager, starting from version 4.17, a new component “VPN Client Interface” is introduced in Hub & Spoke Topology of type (Easy VPN).

Note Some of the policies used in Easy VPN topologies are similar to those used in remote access VPNs. In remote access VPNs, policies are configured between servers and mobile remote PCs running VPN client software, whereas, in site-to-site Easy VPN topologies, the clients are hardware devices.

This section contains the following topics:

- Easy VPN with Dial Backup, page 28-2
- Easy VPN with High Availability, page 28-3
- Easy VPN with Dynamic Virtual Tunnel Interfaces, page 28-3
- Easy VPN Configuration Modes, page 28-3
- Easy VPN and IKE Extended Authentication (Xauth), page 28-4
- Overview of Configuring Easy VPN, page 28-5
- Important Notes About Easy VPN Configuration, page 28-7

Easy VPN with Dial Backup

Dial backup for Easy VPN allows you to configure a dial backup tunnel connection on your remote client device. The backup feature is activated only when real traffic is ready to be sent, eliminating the need for expensive dialup or ISDN links that must be created and maintained even when there is no traffic.

Note Easy VPN dial backup can be configured only on remote clients that are routers running IOS version 12.3(14)T or later.

In an Easy VPN configuration, when a remote device attempts to connect to the server and the tracked IP is no longer accessible, the primary connection is torn down and a new connection is established over the Easy VPN backup tunnel to the server. If the primary hub cannot be reached, the primary configuration switches to the failover hub with the same primary configuration and not to the backup configuration.

Only one backup configuration is supported for each primary Easy VPN configuration. Each inside interface must specify the primary and backup Easy VPN configuration. IP static route tracking must be configured for dial backup to work on an Easy VPN remote device. The object tracking configuration is independent of the Easy VPN remote dial backup configuration. The object tracking details are specified in the spoke’s Edit Endpoints dialog box.

For more information about dial backup, see Configuring Dial Backup, page 25-40.

Easy VPN with High Availability

You can configure High Availability (HA) on devices in an Easy VPN topology. High Availability provides automatic device backup when configured on Cisco IOS routers or Catalyst 6500/7600 devices that run IP over LANs. You can create an HA group made up of two or more hub devices in your Easy VPN that use Hot Standby Routing Protocol (HSRP) to provide transparent, automatic device failover. For more information, see Configuring High Availability in Your VPN Topology, page 25-52.

Easy VPN with Dynamic Virtual Tunnel Interfaces

The IPsec virtual tunnel interface (VTI) feature simplifies the configuration of GRE tunnels that need to be protected by IPsec for remote access links. A VTI is an interface that supports IPsec tunneling, and allows you to apply interface commands directly to the IPsec tunnels. The configuration of a virtual tunnel interface reduces overhead as it does not require a static mapping of IPsec sessions to a particular physical interface where the crypto map is applied.

IPsec VTIs support both unicast and multicast encrypted traffic on any physical interface, such as in the case of multiple paths. Traffic is encrypted or decrypted when it is forwarded from or to the tunnel interface and is managed by the IP routing table. Dynamic or static IP routing can be used to route the traffic to the virtual interface. Using IP routing to forward traffic to the tunnel interface simplifies IPsec VPN configuration compared to the more complex process of using access control lists (ACLs) with a crypto map. Dynamic VTIs function like any other real interface so that you can apply quality of service (QoS), firewall, and other security services as soon as the tunnel is active.

Dynamic VTIs use a virtual template infrastructure for dynamic instantiation and management of IPsec interfaces. In an Easy VPN topology, Security Manager implicitly creates the virtual template interface for the device. If the device is a hub, the user must provide the IP address on the hub that will be used as the virtual template interface—this can be a subnet (pool of addresses) or an existing loopback or physical interface. On a spoke, the virtual template interface is created without an IP address.

In Security Manager, you configure Dynamic VTI in the Easy VPN IPsec Proposal page. See Configuring Dynamic VTI for Easy VPN, page 28-12.

Notes

- Dynamic VTI can be configured only in a hub-and-spoke Easy VPN topology on routers running IOS version 12.4(2)T and later, except 7600 devices. It is not supported on PIX Firewalls, ASA devices, or Catalyst 6000 series switches.
- Not all the hubs/spokes require Dynamic VTI configuration during discovery or provision. You can extend the existing Easy VPN topology (including routers not supporting dVTI) to add routers that support dVTI.
- Dynamic VTI is supported on only servers, only clients (if server does not support dVTI), or both clients and servers.
- You cannot configure High Availability on hubs/servers that have been configured with dVTI.
- You can also configure Dynamic VTI in remote access VPNs. For more information, see Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices), page 33-7.

Easy VPN Configuration Modes

Easy VPN can be configured in three modes—Client, Network Extension, and Network Extension Plus.

- Client mode—The default configuration that allows devices at the client site to access resources at the central site, but disallows access to the central site for resources at the client site. In client mode, a single IP address is pushed to the remote client from the server when the VPN connection is established. This address is typically a routable address in the private address space of the customer network. All traffic passing across the Easy VPN tunnel undergoes Port Address Translation (PAT) to that single pushed IP address.
- Network Extension mode—Allows users at the central site to access the network resources at the client site, and allows the client PCs and hosts direct access to the PCs and hosts at the central site. Network Extension mode specifies that the hosts at the client end of the VPN tunnel should be given IP addresses that are fully routable and reachable by the destination network. The devices at both ends of the connection will form one logical network. PAT is not used, so the hosts at the client end have direct access to the hosts at the destination network. In other words, the Easy VPN server (the hub) gives routable addresses to the Easy VPN client (the spoke), while the whole LAN behind the client will not undergo PAT.
- Network Extension Plus mode—An enhancement to Network Extension mode, which can be configured only on IOS routers. It enables an IP address that is received via mode configuration to be automatically assigned to an available loopback interface. This IP address can be used for connecting to your router for remote management and troubleshooting (ping, Telnet, and Secure Shell). If you select this option an some clients are not IOS routers, those clients are configured in Network Extension mode.

Note All modes of operation can also support split tunneling, which allows secure access to corporate resources through the VPN tunnel while also allowing Internet access through a connection to an ISP or other service (thereby eliminating the corporate network from the path for web access).

You configure the mode in the Client Connection Characteristics policy as described in Configuring Client Connection Characteristics for Easy VPN, page 28-7.

Related Topics

- Important Notes About Easy VPN Configuration, page 28-7
- Understanding Easy VPN, page 28-1

Easy VPN and IKE Extended Authentication (Xauth)

When negotiating tunnel parameters for establishing IPsec tunnels in an Easy VPN configuration, IKE Extended Authentication (Xauth) adds another level of authentication that identifies the user who requests the IPsec connection. If the VPN server is configured for Xauth, the client waits for a username/password challenge after the IKE security association (SA) has been established. When the end user responds to the challenge, the response is forwarded to the IPsec peers for an additional level of authentication.

The information that is entered is checked against authentication entities using authentication, authorization, and accounting (AAA) protocols such as RADIUS and TACACS+. Token cards may also be used via AAA proxy. During Xauth, a user-specific attribute can be retrieved if the credentials of that user are validated via RADIUS.

Note VPN servers that are configured to handle remote clients should always be configured to enforce user authentication.

Security Manager allows you to save the Xauth username and password on the device itself so you do not need to enter these credentials manually each time the Easy VPN tunnel is established. The information is saved in the device's configuration file and used each time the tunnel is established. Saving the credentials in the device's configuration file is typically used if the device is shared between several PCs and you want to keep the VPN tunnel up all the time, or if you want the device to automatically bring up the tunnel whenever there is traffic to be sent.

Saving the credentials in the device's configuration file, however, could create a security risk, because anyone who has access to the device configuration can obtain this information. An alternative method for Xauth authentication is to manually enter the username and password each time Xauth is requested. You can select whether to use a web browser window or the router console to enter the credentials. Using web-based interaction, a login page is returned, in which you can enter the credentials to authenticate the VPN tunnel. After the VPN tunnel comes up, all users behind this remote site can access the corporate LAN without being prompted again for the username and password. Alternatively, you can choose to bypass the VPN tunnel and connect only to the Internet, in which case a password is not required.

Easy VPN Tunnel Activation

If the device credentials (Xauth username and password) are stored on the device itself, you must select a tunnel activation method for IOS router clients. Two options are available:

- Auto—The Easy VPN tunnel is established automatically when the Easy VPN configuration is delivered to the device configuration file. If the tunnel times out or fails, the tunnel automatically reconnects and retries indefinitely. This is the default option.
- Traffic Triggered Activation—The Easy VPN tunnel is established whenever outbound local (LAN side) traffic is detected. Traffic Triggered Activation is recommended for use with the Easy VPN dial backup configuration so that backup is activated only when there is traffic to send across the tunnel. When using this option, you must specify the Access Control List (ACL) that defines the "interesting" traffic.

Note Manual tunnel activation is configured implicitly if you select to configure the Xauth password interactively. In this case, the device waits for a command before attempting to establish the Easy VPN remote connection. When the tunnel times out or fails, subsequent connections will also have to wait for the command.

You configure the xauth and tunnel activation mode in the Client Connection Characteristics policy as described in Configuring Client Connection Characteristics for Easy VPN, page 28-7.

Related Topics

- Important Notes About Easy VPN Configuration, page 28-7
- Understanding Easy VPN, page 28-1
- Configuring Credentials Policy Objects, page 28-9

Overview of Configuring Easy VPN

When a remote client initiates a connection to a VPN server, device authentication between the peers occurs using IKE, followed by user authentication using IKE Extended Authentication (Xauth), VPN policy push (in Client, Network Extension, or Network Extension Plus mode), and IPsec security association (SA) creation.

The following provides an overview of this process:

1. The client initiates IKE Phase 1 via aggressive mode if a preshared key is to be used for authentication, or main mode if digital certificates are used. If the client identifies itself with a preshared key, the accompanying user group name (defined during configuration) is used to identify the group profile associated with this client. If digital certificates are used, the organizational unit (OU) field of a distinguished name (DN) is used to identify the user group name. See PKI Enrollment Dialog Box—Certificate Subject Name Tab, page 26-66.

   Note Because the client may be configured for preshared key authentication, which initiates IKE aggressive mode, the administrator should change the identity of the VPN device via the crypto isakmp identity hostname command. This will not affect certificate authentication via IKE main mode.

2. The client attempts to establish an IKE SA between its public IP address and the public IP address of the VPN server. To reduce the amount of manual configuration on the client, every combination of encryption and hash algorithms, in addition to authentication methods and D-H group sizes, is proposed.
3. Depending on its IKE policy configuration, the VPN server determines which proposal is acceptable to continue negotiating Phase 1.

   Note Device authentication ends and user authentication begins at this point.

4. After the IKE SA is successfully established, and if the VPN server is configured for Xauth, the client waits for a “username/password” challenge and then responds to the challenge of the peer. The information that is entered is checked against authentication entities using authentication, authorization, and accounting (AAA) protocols such as RADIUS and TACACS+. Token cards may also be used via AAA proxy. During Xauth, a user-specific attribute can be retrieved if the credentials of that user are validated via RADIUS.

   Note VPN servers that are configured to handle remote clients should always be configured to enforce user authentication.

5. If the server indicates that authentication was successful, the client requests further configuration parameters from the peer. The remaining system parameters (for example, IP address, DNS, and split tunnel attributes) are pushed to the client using client or network extension mode configuration.

   Note The IP address pool and group preshared key (if Rivest, Shamir, and Adelman [RSA] signatures are not being used) are the only required parameter in a group profile. All other parameters are optional.

6. After each client is assigned an internal IP address via mode configuration, Reverse Route Injection (RRI), if configured, ensures that a static route is created on the device for each client internal IP address.
7. IKE quick mode is initiated to negotiate and create IPsec SAs.

The connection is complete.

Important Notes About Easy VPN Configuration

Before you configure an Easy VPN policy in your topology, you should know the following:

- In an Easy VPN topology configuration, deployment fails if a 72xx series router is used as a remote client device. The Easy VPN client is supported on PIX 501, 506, 506E Firewalls running PIX 6.3, Cisco 800-3900 Series routers, and ASA 5505 devices running ASA Software release 7.2 or later.
- If you try to configure a Public Key Infrastructure (PKI) policy on a PIX 6.3 remote client in an Easy VPN topology configuration, deployment fails. For successful deployment on this device, you must first issue the PKI certificate on the CA server, and then try again to deploy the device. For more information about PKI policies, see Understanding Public Key Infrastructure Policies, page 26-51.
- In some cases, deployment fails on a device that serves as an Easy VPN client if the crypto map is configured on the NAT (or PAT) internal interface instead of the external interface. On some platforms, the inside and outside interfaces are fixed. For example, on a Cisco 1700 series router the VPN interface must be the device's FastEthernet0 interface. On a Cisco 800 series router the VPN interface could be either the device's Ethernet0 or Dialer1 interface, depending on the configuration. On a Cisco uBR905/uBR925 cable access router, the VPN interface must be the Ethernet0 interface.

Configuring Client Connection Characteristics for Easy VPN

Use the Client Connection Characteristics page to specify how traffic will be routed in the Easy VPN topology and how the VPN tunnel will be established. The characteristics defined in this policy are configured on the remote clients. Before configuring this policy, read the following topics:

- Easy VPN Configuration Modes, page 28-3
- Easy VPN and IKE Extended Authentication (Xauth), page 28-4

Navigation Path

- (Site-to-Site VPN Manager Window, page 25-18) Select an Easy VPN topology in the VPNs selector, then select **Client Connection Characteristics** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > Client Connection Characteristics** and create a new policy or edit an existing policy.

Related Topics

- Understanding Easy VPN, page 28-1
- Creating Access Control List Objects, page 6-53
- Important Notes About Easy VPN Configuration, page 28-7

Field Reference

Table 28-1 Easy VPN Client Connection Characteristics Page

Element	Description
Mode	The configuration mode for the remote devices: • **Client**—Specifies that all traffic from the remote client's inside network will undergo Port Address Translation (PAT) to a single IP address which was assigned for the device by the head end server at connect time. • **Network Extension**—Specifies that PCs and other hosts at the client end of the VPN tunnel should be given IP addresses that are fully routable and reachable by destination network. PAT is not used, allowing the client PCs and hosts to have direct access to the PCs and hosts at the destination network. • **Network Extension Plus**—An enhancement to Network Extension mode, that enables an IP address that is received via mode configuration to be automatically assigned to an available loopback interface. The IPsec SAs for this IP address are automatically created by the Easy VPN client. The IP address is typically used for troubleshooting (using ping, Telnet, and Secure Shell). If you select Network Extension Plus, this mode is configured on IOS routers only. Clients that are PIX or ASA devices are configured in Network Extension mode. For more information, see Easy VPN Configuration Modes, page 28-3.
Xauth Credentials Source	Select how you want to enter the Xauth credentials for user authentication when you establish a VPN connection with the server: • Device Stored Credentials (default)—The username and password are saved on the device itself in the device's configuration file to be used each time the tunnel is established. • Interactive Entered Credentials—Enables you to manually enter the username and password each time Xauth is requested, in a web browser window or from the router console. For more information, see Easy VPN and IKE Extended Authentication (Xauth), page 28-4.
Xauth Credentials	Available only if you selected **Device Stored Credentials** as the Xauth Credentials Source. The credentials policy object that defines the default Xauth credentials. Enter the name of the object or click **Select** to select it from a list or to create a new object. For more information, see Configuring Credentials Policy Objects, page 28-9. **Note** If you want to configure different Xauth credentials on your remote client, you must configure the credentials policy object to allow overrides (select **Allow Value Override per Device** in the object definition).

Table 28-1 Easy VPN Client Connection Characteristics Page (continued)

Element	Description
Tunnel Activation (IOS)	Available only if you selected the **Device Stored Credentials** option for the Xauth password source. For IOS router clients, select a tunnel activation method: • **Auto** (default)—The Easy VPN tunnel is established automatically when the Easy VPN configuration is delivered to the device configuration file. If the tunnel times out or fails, the tunnel automatically reconnects and retries indefinitely. • **Traffic Triggered Activation**—The Easy VPN tunnel is established whenever outbound local (LAN side) traffic is detected. If you select traffic triggered activation, also enter the name of the Access Control List (ACL) policy object that defines the traffic that should activate the tunnel. Click **Select** to select the object or to create a new object. Traffic Triggered Activation is recommended for use when Easy VPN dial backup is configured so that backup is activated only when there is traffic to send across the tunnel. **Note** Manual tunnel activation is configured implicitly when you select to configure the Xauth password interactively.
User Authentication Method (IOS)	Available only if you selected the **Interactive Entered Credentials** option for the Xauth credentials source. The option applies to remote IOS routers only. Select one of these ways to enter the Xauth username and password interactively each time Xauth authentication is requested: • Web Browser (default)—Manually in a web browser window. • Router Console—Manually from the router's command line.

Configuring Credentials Policy Objects

Use the Credentials dialog box to create, copy and edit Credential objects.

Credential objects are used in Easy VPN configuration during IKE Extended Authentication (Xauth) when authenticating user access to the network and network services. When negotiating tunnel parameters for establishing IPsec tunnels in an Easy VPN configuration, Xauth identifies the user who requests the IPsec connection. If the VPN server is configured for Xauth, the client waits for a "username/password" challenge after the IKE SA has been established. When the end user responds to the challenge, the response is forwarded to the IPsec peers for an additional level of authentication. You can save the Xauth credentials (username and password) on the device itself so you do not need to enter them manually each time the Easy VPN tunnel is established.

Navigation Path

Select **Manage > Policy Objects**, then select **Credentials** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Easy VPN and IKE Extended Authentication (Xauth), page 28-4

- Configuring Client Connection Characteristics for Easy VPN, page 28-7
- Policy Object Manager, page 6-4

Field Reference

Table 28-2 Credentials Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object (up to 1024 characters).
Username	The name that will be used to identify the user during Xauth authentication.
Password Confirm	The password for the user, entered in both fields. The password must be alphanumeric and a maximum of 128 characters. Spaces are not allowed.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring an IPsec Proposal for Easy VPN

Use the Easy VPN IPsec Proposal page to configure the IPsec proposal used during IKE Phase 2 negotiations for Easy VPN topologies. The IPsec proposal is configured on the IPSec Proposal tab; the options are described below.

In Easy VPN topologies, you can also configure a dynamic virtual interface on the Dynamic VTI tab. For an explanation of dVTI configuration, see Configuring Dynamic VTI for Easy VPN, page 28-12.

Note This topic describes the IPsec Proposal page when the site-to-site VPN technology is Easy VPN. For a description of the IPsec Proposal page when the site-to-site VPN technology is something else, see Configuring IPsec Proposals in Site-to-Site VPNs, page 26-22.

Navigation Path

- (Site-to-Site VPN Manager Window, page 25-18) Select an Easy VPN topology in the VPNs selector, then select **Easy VPN IPsec Proposal** in the Policies selector. Click the **IPSec Proposal** tab.
- (Policy view) Select **Site-to-Site VPN > Easy VPN IPsec Proposal** from the Policy Types selector. Select an existing shared policy or create a new one. Click the **IPSec Proposal** tab.

Related Topics

- Understanding Easy VPN, page 28-1
- Configuring an IPsec Proposal for Easy VPN, page 28-10
- Understanding AAA Server and Server Group Objects, page 6-27
- Understanding IPsec Proposals, page 26-18

Field Reference

Table 28-3 Easy VPN IPsec Proposal Tab

Element	Description
IKEv1 Transform Sets	The transform sets to be used for your tunnel policy. Transform sets specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel. You can select up to 11 transform sets. For more information, see Understanding Transform Sets, page 26-20. Transform sets may use only tunnel mode IPsec operation. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security will be used. Click **Select** to select the IPsec transform set policy objects to use in the topology. If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27.
Reverse Route	Supported on ASA 5500 series devices, PIX 7.0+ devices, and Cisco IOS routers except 7600 devices. Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint. For more information, see Understanding Reverse Route Injection, page 26-21. Select one of the following options to configure RRI on the crypto map: • **None**—Disables the configuration of RRI on the crypto map. • **Standard**—(ASA, PIX 7.0+, IOS devices) Creates routes based on the destination information defined in the crypto map access control list (ACL). This is the default option. • **Remote Peer**—(IOS devices only) Creates two routes, one for the remote endpoint and one for route recursion to the remote endpoint via the interface to which the crypto map is applied. • **Remote Peer IP**—(IOS devices only) Specifies an address as the explicit next hop to the remote VPN device. Enter the IP address or a network/host object that specifies the address, or click **Select** to select the network/host object from a list or to create a new object. **Note** If you use network/host objects, you can select the **Allow Value Override per Device** option in the object to override the IP address, if required, for specific devices that use this object.

Table 28-3 Easy VPN IPsec Proposal Tab (continued)

Element	Description
Enable Network Address Translation Traversal	Supported on PIX 7.0+ and ASA 5500 series devices. Whether to allow Network Address Translation (NAT) traversal. Use NAT traversal when there is a device between a VPN-connected hub and spoke, and that performs Network Address Translation (NAT) on the IPsec traffic. For information about NAT traversal, see Understanding NAT in VPNs, page 26-41
Group Policy Lookup/AAA Authorization Method	Supported on Cisco IOS routers only. The AAA authorization method list that will be used to define the order in which the group policies are searched. Group policies can be configured on both the local server or on an external AAA server. Remote users are grouped, so that when the remote client establishes a successful connection to the VPN server, the group policies for that particular user group are pushed to all clients belonging to the user group. Click **Select** to open a dialog box that lists all available AAA group servers, and in which you can create AAA group server objects. Select all that apply and use the up and down arrow buttons to put them in priority order.
User Authentication (Xauth)/AAA Authentication Method	Supported on Cisco IOS routers and PIX 6.3 firewalls only. The AAA or Xauth user authentication method used to define the order in which user accounts are searched. Xauth allows all AAA authentication methods to perform user authentication in a separate phase after the IKE authentication phase 1 exchange. The AAA configuration list-name must match the Xauth configuration list-name for user authentication to occur. After the IKE SA is successfully established, and if the device is configured for Xauth, the client waits for a username/password challenge and then responds to the challenge of the peer. The information that is entered is checked against authentication entities using authentication, authorization, and accounting (AAA) protocols such as RADIUS and TACACS+. Click **Select** to open a dialog box that lists all available AAA group servers, and in which you can create AAA group server objects. Select all that apply and use the up and down arrow buttons to put them in priority order.

Configuring Dynamic VTI for Easy VPN

Use the Dynamic VTI tab of the Easy VPN IPSec Proposal policy to configure a dynamic virtual tunnel interface on a device in a hub-and-spoke Easy VPN topology. For more information, see Easy VPN with Dynamic Virtual Tunnel Interfaces, page 28-3.

Note Dynamic VTI can be configured only on IOS routers running IOS version 12.4(2)T and later, except 7600 devices.

Navigation Path

- (Site-to-Site VPN Manager Window, page 25-18) Select an Easy VPN topology in the VPNs selector, then select **Easy VPN IPsec Proposal** in the Policies selector. Click the **Dynamic VTI** tab.
- (Policy view) Select **Site-to-Site VPN > Easy VPN IPsec Proposal** from the Policy Types selector. Select an existing shared policy or create a new one. Click the **Dynamic VTI** tab.

Related Topics

- Understanding Easy VPN, page 28-1
- Configuring an IPsec Proposal for Easy VPN, page 28-10

Field Reference

Table 28-4 Easy VPN IPSec Proposal, Dynamic VTI Tab

Element	Description
Enable Dynamic VTI	When selected, enables Security Manager to implicitly create a dynamic virtual template interface on the device. If the device is a hub server that does not support Dynamic VTI, a warning message is displayed, and a crypto map is deployed without dynamic VTI. In the case of a client device, an error message is displayed.
Virtual Template IP	If you are configuring Dynamic VTI on a hub in the topology, specify either the subnet address or interface role: • **Subnet**—To use the IP address taken from a pool of addresses. Enter the private IP address including the subnet mask, for example 10.1.1.0/24. • **Interface Role**—To use a physical or loopback interface on the device. If required, click **Select** to open the Interface selector where you can select the interface role object that identifies the desired interface. If an appropriate object does not already exist, you can create one in the selection dialog box. If you are configuring Dynamic VTI on a spoke in the topology, select **None**.

Configuring a Connection Profile Policy for Easy VPN

A connection profile consists of a set of records that contain IPsec tunnel connection policies. Connection profiles, or tunnel groups, identify the group policy for a specific connection, and include user-oriented attributes. If you do not assign a particular group policy to a user, the default group policy for the connection applies. For a successful connection, the username of the remote client must exist in the database, otherwise the connection is denied.

In site-to site VPNs, you configure connection profile policies on an Easy VPN server, which can be a PIX Firewall version 7.0+ or an ASA 5500 series device. The Easy VPN connection profile policy is similar to the one used for remote access VPNs.You can unassign the connection profile policy if none of the Easy VPN servers are ASA or PIX 7.0+ devices.

Creating a connection profile policy involves specifying:

- The group policy—A collection of user-oriented attributes stored either internally on the device or externally on RADIUS/LDAP server.
- Global AAA settings—Authentication, Authorization, and Accounting servers.
- The DHCP servers to be used for client address assignment, and the address pools from which the IP addresses will be assigned.
- Settings for Internet Key Exchange (IKE) and IPsec (such as preshared key).

On the PIX7.0+/ASA Connection Profiles page, you can connection profiles on your Easy VPN server.

Related Topics

- Creating or Editing VPN Topologies, page 25-28
- Understanding IPsec Technologies and Policies, page 25-5
- Understanding Easy VPN, page 28-1

Step 1 Do one of the following:

- (Site-to-Site VPN Manager Window, page 25-18) Select an Easy VPN topology in the VPNs selector, then select **Connection Profiles (PIX 7.0/ASA)** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > Connection Profiles (PIX 7.0/ASA)** from the Policy Types selector. Select an existing shared policy or create a new one.

For information on the policy, see Connection Profiles Page, page 31-8.

Step 2 On the **General** tab, specify the connection profile name and group policies and select which method (or methods) of address assignment to use. For a description of the available properties, see General Tab (Connection Profiles), page 31-10.

Step 3 Click the **AAA** tab and specify the AAA authentication parameters for an the connection profile. For a description of the elements on the tab, see AAA Tab (Connection Profiles), page 31-13.

Step 4 Click the **IPsec** tab and specify IPsec and IKE parameters for the connection profile. For a description of the elements on the tab, see IPSec Tab (Connection Profiles), page 31-19.

Configuring a User Group Policy for Easy VPN

Use the User Group Policy page to create or edit a user group policy on your Easy VPN server. When you configure an Easy VPN server, you create a user group to which remote clients belong. An Easy VPN user group policy can be configured on a Cisco IOS security router, PIX 6.3 Firewall, or Catalyst 6500 /7600 device. You can unassign the user group policy if none of the Easy VPN servers are IOS routers, Catalyst 6500/7600 devices, or PIX 6.3 firewalls.

Remote clients must have the same group name as the user group configured on the server in order to connect to the device, otherwise no connection is established. When the remote client establishes a successful connection to the VPN server, the group policies for that particular user group are pushed to all clients belonging to the user group.

Select the user group policy object that you want to use in the policy from the Available User Groups list. You can create a new user group object by clicking the **Create (+)** button, or edit an existing group by selecting it and clicking the **Edit (pencil icon)** button. For information about configuring the user group object, see Add or Edit User Group Dialog Box, page 34-73.

Navigation Path

- (Site-to-Site VPN Manager Window, page 25-18) Select an Easy VPN topology in the VPNs selector, then select **User Group Policy** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > User Group Policy** from the Policy Types selector. Select an existing shared policy or create a new one.

Related Topics

- Understanding Easy VPN, page 28-1

Group Encrypted Transport (GET) VPNs

Cisco Group Encrypted Transport virtual private network (GET VPN) is a full-mesh VPN technology that can be used in a variety of WAN environments, including IP and Multiprotocol Label Switching (MPLS). GET VPN comprises a set of features that are necessary to secure IP multicast group traffic or unicast traffic over a private WAN that originates on or flows through a Cisco IOS device. GET VPN combines the keying protocol Group Domain of Interpretation (GDOI) with IP security (IPsec) encryption to provide users with an efficient method to secure IP multicast or unicast traffic. GET VPN enables the router to apply encryption to nontunneled (that is, "native") IP multicast and unicast packets and eliminates the requirement to configure tunnels to protect multicast and unicast traffic.

Cisco Group Encrypted Transport VPN provides the following benefits:

- Provides data security and transport authentication, helping to meet security compliance and internal regulation by encrypting all WAN traffic.
- Enables high-scale network meshes and eliminates complex peer-to-peer key management with group encryption keys.
- For Multiprotocol Label Switching (MPLS) networks, maintains network intelligence such as full-mesh connectivity, natural routing path, and Quality of Service (QoS).
- Grants easy membership control with a centralized key server.
- Helps ensure low latency and jitter by enabling full-time, direct communications between sites, without requiring transport through a central hub.
- Reduces traffic loads on customer premises equipment (CPE) and provider-edge (PE) encryption devices by using the core network for replication of multicast traffic, avoiding packet replication at each individual peer site.

Tip For information about the CLI configuration of GET VPN, see *Cisco Group Encrypted Transport VPN* on Cisco.com.

This chapter contains the following topics:

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Understanding the GET VPN Registration Process, page 29-4
- Understanding the GET VPN Security Policy and Security Associations, page 29-10
- Configuring GET VPN, page 29-12
- Generating and Synchronizing RSA Keys, page 29-13
- Configuring the IKE Proposal for GET VPN, page 29-15

- Configuring Global Settings for GET VPN, page 29-16
- Configuring GET VPN Key Servers, page 29-18
- Configuring GET VPN Group Members, page 29-20
- Using Passive Mode to Migrate to GET VPN, page 29-23
- Troubleshooting GET VPN Configurations, page 29-25

Understanding Group Encrypted Transport (GET) VPNs

Networked applications such as voice and video increase the need for instantaneous, branch-interconnected, and QoS-enabled WANs. The distributed nature of these applications results in increased demands for scale. At the same time, enterprise WAN technologies force businesses to trade off between QoS-enabled branch interconnectivity and transport security. As network security risks increase and regulatory compliance becomes essential, Group Encrypted Transport VPN (GET VPN), a WAN encryption technology, eliminates the need to compromise between network intelligence and data privacy.

With GET, Cisco provides tunnelless VPN, which eliminates the need for IPsec tunnels. By removing the need for point-to-point tunnels, meshed networks can scale higher while maintaining network-intelligence features critical to voice and video quality. GET is a standards-based security model that is based on the concept of a trusted group to eliminate point-to-point IPsec tunnels and their associated overlay routing. Trusted group members share a common security association (SA), also known as a group SA. This enables group members to decrypt traffic that was encrypted by any other group member. By using trusted groups instead of point-to-point tunnels, full-mesh networks can scale higher while maintaining network-intelligence features (such as QoS, routing, and multicast), which are critical to voice and video quality.

GET-based networks can be used in a variety of WAN environments, including IP and Multiprotocol Label Switching (MPLS). MPLS VPNs that use this encryption technology are highly scalable, manageable, and cost-effective, and they meet government-mandated encryption requirements. The flexible nature of GET allows security-conscious enterprises either to manage their own network security over a service provider WAN service or to offload encryption services to their providers. GET simplifies securing large Layer 2 or MPLS networks that require partial or full-mesh connectivity.

In addition to leveraging the existing IKE, IPsec and multicast technologies, a GET VPN topology includes these key elements and features:

- Group members—The routers that exchange the actual traffic within the VPN are called group members. Group members provide encryption services to the traffic. Encryption policies are defined centrally on the key server and downloaded to the group member at the time of registration. Based on these downloaded policies, a group member decides whether traffic needs to be encrypted or decrypted and what keys to use.

  Although group members primarily obtain encryption policies from the key server, you can configure local service policy ACLs on the group members to exclude traffic from encryption based on local requirements. For more information, see Understanding the GET VPN Security Policy and Security Associations, page 29-10.

  Note A device can be a group member of more than one group.

- Key servers—The routers that act as key servers are the gatekeepers to the topology. The group member must successfully register with a key server before becoming an active member of the VPN. The key servers control the shared service policy, and generate and transmit keys to group members. Key servers cannot be group members themselves, but a single key server can service more than one topology. For more information, see Understanding the GET VPN Registration Process, page 29-4.
- The Group Domain of Interpretation (GDOI) group key management protocol is used to provide a set of cryptographic keys and policies to a group of devices. In a GET VPN network, GDOI is used to distribute common IPsec keys to a group of enterprise VPN gateways (group members) that must communicate securely. Devices designated as key servers periodically refresh and send out the updated keys to the group members using a process called "rekeying."

  The GDOI protocol uses the Phase 1 Internet Key Exchange (IKE) SA. All participating VPN gateways authenticate themselves to the device providing keys using IKE. All IKE authentication methods, for example, pre-shared keys (PSKs) and public key infrastructure (PKI), are supported for initial authentication. After the VPN gateways are authenticated and provided with the appropriate security keys using the IKE SA, the IKE SA expires and GDOI is used to update the group members in a more scalable and efficient manner. For more information about GDOI, refer to RFC 3547.
- Address preservation—IPsec-protected data packets carry the original source and destination in the outer IP header rather than replacing them with tunnel endpoint addresses. Address preservation allows GET VPN to use the routing functionality present within the core network. Address preservation allows routing to deliver the packets to any customer-edge (CE) device in the network that advertises a route to the destination address. Any source and destination matching the policy for the group will be treated in a similar manner. In the situation where a link between IPsec peers is not available, address preservation also helps combat traffic "black-hole" situations.

  Header preservation also maintains routing continuity throughout the enterprise address space and in the WAN. As a result, end host addresses of the campus are exposed in the WAN (for MPLS, this applies to the edge of the WAN). For this reason, GET VPN is applicable only when the WAN network acts as a "private" network (for example, in an MPLS network).

The following figure shows the general operation of a GET VPN topology.

Figure 29-1 General GET VPN Operation

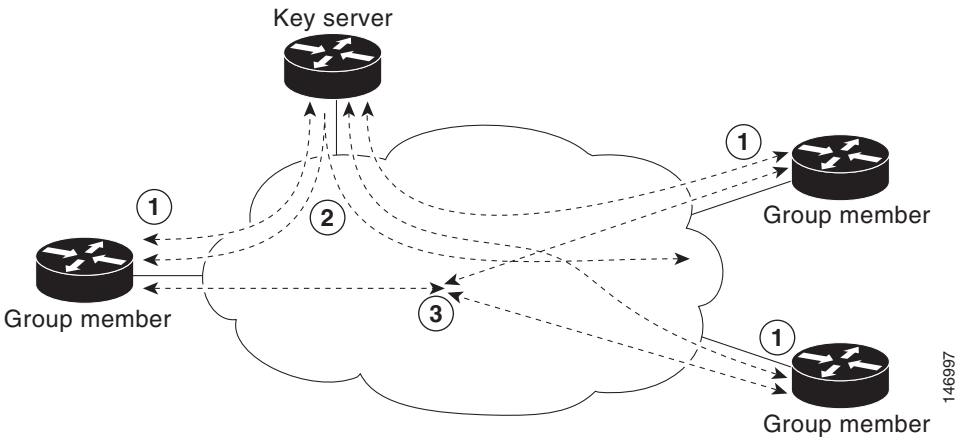

1. Group members register with the key server using the Group Domain of Interpretation (GDOI) protocol. The key server authenticates and authorizes the group members and downloads the IPsec policy and keys that are necessary for them to encrypt and decrypt IP multicast and unicast packets. The registration process can use unicast or multicast communications.

2. Group members exchange IP packets that are encrypted using IPsec. Only the group members are an active part of the VPN.
3. As needed, the key server pushes a rekey message to the group members. The rekey message contains new IPsec policy and keys to use when old IPsec security associations (SAs) expire. Rekey messages are sent in advance of the SA expiration time to ensure that valid group keys are always available.

GET VPN is provisioned using Security Manager with the following caveats:

- GET VPN-aware VRF is not supported.
- DMVPN with GET is not supported, because there is no way to define DMVPN without tunnel protection in Security Manager.
- Manual configuration of a group member to join a multicast group (ip igmp join-group) is not supported. Security Manager only provisions static source-specific multicast (SSM) mappings.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Understanding the GET VPN Security Policy and Security Associations, page 29-10
- Configuring GET VPN, page 29-12

Understanding the GET VPN Registration Process

In GET VPN, group members comprise the VPN topology. Traffic in the VPN is traffic between group members. For a device to become a group member, the device must successfully register with a key server. Key servers maintain the security association (SA) policy and create and maintain the keys for the group. When a group member registers, the key server downloads the policy and the keys to the group member. The key server also rekeys the group before existing keys expire.

The key server has two responsibilities: servicing registration requests and sending rekeys. A group member can register at any time and receive the most current policy and keys. When a group member registers with the key server, the key server verifies the group ID that the group member is attempting to join. If the group ID is valid, the key server sends the security association policy to the group member. After the group member acknowledges that it can handle the downloaded policy, the key server downloads the respective keys.

Communication among the key server and group members is encrypted and secured using two types of keys: the traffic encryption key (TEK) and the key encryption key (KEK). The TEK is downloaded by the key server to all the group members. The downloaded TEK is used by all the group members to communicate securely among each other. This key is essentially the group key that is shared by all the group members. The group policies and IPsec SAs are refreshed by the key server using periodic rekey messages to the group members. The KEK is also downloaded by the key server and is used by the group members to decrypt the incoming rekey messages from the key server.

The key server sends out rekey messages either because of an impending IPsec SA expiration or because the security policy has changed on the key server. A rekey can also happen if the KEK timer has expired (the key server sends out a KEK rekey). Rekey messages might also be retransmitted periodically to account for possible packet loss. If the rekey mechanism is multicast, there is no efficient feedback mechanism by which receivers can indicate that they did not receive a rekey message, so retransmission seeks to bring all receivers up to date. If the rekey mechanism is unicast, the receivers send an acknowledgment message.

The key server generates the group policy and IPsec security associations (SAs) for the GDOI group. The information generated by the key server includes multiple TEK attributes, traffic encryption policy, lifetime, source and destination, a Security Parameter Index (SPI) ID that is associated with each TEK, and the rekey policy (one KEK). Note that the group member might also have a local security policy configured that is merged with the one downloaded; for complete information see Understanding the GET VPN Security Policy and Security Associations, page 29-10.

The following figure illustrates the communication flow between group members and the key server. The key server, after receiving registration messages from a group member, generates the information that contains the group policy and new IPsec SAs. The new IPsec SA is then downloaded to the group member. The key server maintains a table that contains the IP address of each group member per group. When a group member registers, the key server adds its IP address in its associated group table, thus allowing the key server to monitor an active group member. A key server can support multiple groups. A group member can be part of multiple groups.

Figure 29-2 Communication Flow Between Group Members and the Key Server

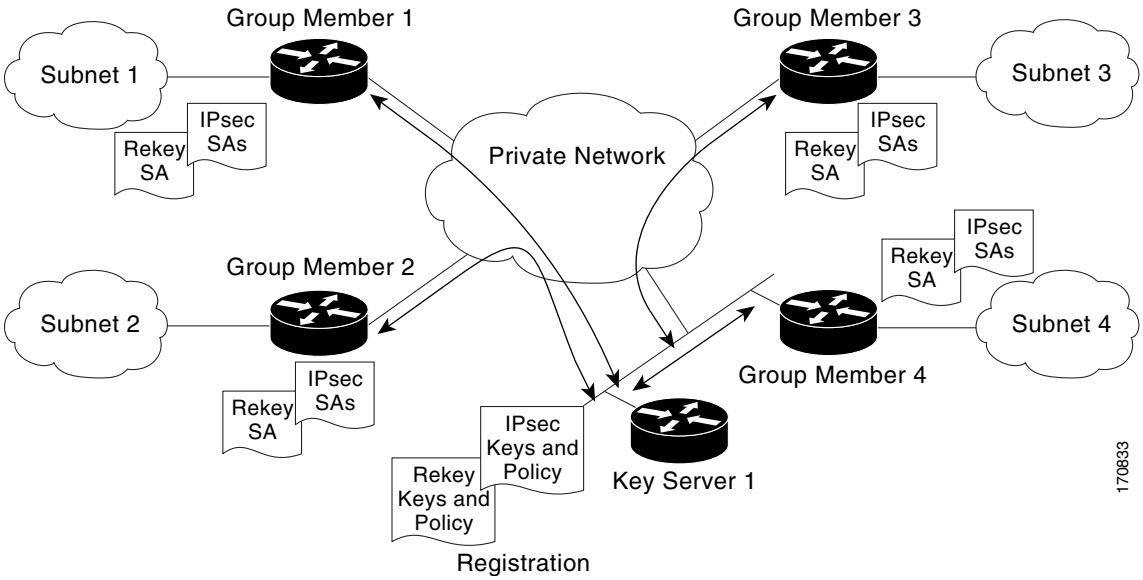

When you configure the GET VPN topology, you can configure the following registration-related features:

- Decide whether to use unicast or multicast for group registration and rekeying. For more information, see Choosing the Rekey Transport Mechanism, page 29-6.

Note If you use multicast, you need to enable multicast on the key servers and group members manually. Security Manager does not provision multicast commands.

- Decide whether to configure more than one key server to provide redundancy and load balancing. For more information, see Configuring Redundancy Using Cooperative Key Servers, page 29-7.
- Decide whether to configure fail-close mode on group members to protect their traffic prior to successful registration with the key server. For more information, see Configuring Fail-Close to Protect Registration Failures, page 29-8.

- Decide whether to require authorization for group members to join the group. You can use certificate authorization (which requires that you also configure the Public Key Infrastructure policy) or preshared keys. Configuring authorization is required if the key server serves more than one group. For information about the configuration options, see the Authorization Type setting described in Defining GET VPN Group Encryption, page 25-54.

Related Topics

- Generating and Synchronizing RSA Keys, page 29-13
- Configuring GET VPN, page 29-12

Choosing the Rekey Transport Mechanism

When you configure the rekey settings in the Group Encryption Policy (as described in Defining GET VPN Group Encryption, page 25-54), you must select whether to use multicast or unicast as the rekey transport mechanism. The key server uses this method whenever sending new keys and IPsec security associations (SAs) to group members or each other. There are advantages and disadvantages to each method.

Multicast is the standard choice. Using multicast, the key server sends one copy of each rekey message to all group members at once using a multicast group address, so there is no rekey delay and group members can install the updated security policy essentially simultaneously (not accounting for regular network delay). However, in some networks, multicast is either an extra cost feature, or it is simply not allowed. If you configure multicast, you must supply the multicast address that will be used by the GET VPN topology.

Unicast can be used when multicast is unavailable or undesirable. Using unicast, the key server sends directed rekey and IPsec SAs to group members, and the group member sends an acknowledgment that the message was received. Because unicast requires sending direct messages and receiving acknowledgments, the key server sends the unicast messages to a subset of the group members at a time (unless you have a relatively small VPN, perhaps fewer than 30 group members, in which case all group members might be sent messages at the same time).

Thus, the relative benefits of multicast and unicast include the following:

- With multicast, the key server does not know if a group member receives a message, whereas with unicast, there are acknowledgments. With unicast, if the key server does not receive the acknowledgment, it resends the message.
- Multicast is faster than unicast, especially for large topologies with hundreds of group members. Multicast rekey uses the same low CPU overhead whether there is one group member in the group or a few thousand.
- With unicast, if a group member continuously fails to send acknowledgments, the key server decides the group member is no longer there and stops sending rekey messages. Thus, the key server always has a list of active group members. The unresponsive group member must reregister to rejoin the GET VPN topology. Because multicast does not use acknowledgments, the key server does not know if a group member becomes unresponsive, and it does not maintain a list of active group members.

Tip To use multicast, you must enable multicast on the key servers and group members. Security Manager does not provision these commands; it only enables multicast rekey, it does not enable the router to send and receive multicast traffic. Therefore, you must manually enable multicast on the device, or use the FlexConfig policy to provision the commands (see Creating FlexConfig Policy Objects, page 7-28).

Fortunately, it is possible to mix multicast and unicast in a single GET VPN topology so long as all key servers support multicast. When deciding which transport mechanism to use, consider the following recommendations:

- If all key servers and group members, and the network, support multicast, use multicast.
- If all of the key servers and most of the group members support multicast, but a small number of group members do not support multicast, use multicast. Group members that do not support multicast will not receive rekey and IPsec SA updates. However, when the lifetime settings for these items are about to expire, unicast group members will reregister with the key server and obtain the new keys and IPsec SAs.
- If no group members, or only a few, support multicast, use unicast. The group members will then receive rekeys and IPsec SA updates from the key server and not need to reregister to get them.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Generating and Synchronizing RSA Keys, page 29-13
- Configuring GET VPN, page 29-12

Configuring Redundancy Using Cooperative Key Servers

The key server is the most important entity in the GET VPN network because the key server maintains the control plane. Therefore, a single key server is a single point of failure for an entire GET VPN network. Because redundancy is an important consideration for key servers, GET VPN supports multiple key servers, called cooperative (COOP) key servers, to ensure seamless fault recovery if a key server fails or becomes unreachable.

You can configure a group member to register to any available key server from a list of all COOP key servers. The group member configuration determines the registration order (see Configuring GET VPN Group Members, page 29-20 and Edit Group Member Dialog Box, page 29-21). The key server defined first is contacted first, followed by the second defined key server, and so on. It is a best practice to distribute group member registration to all available COOP key servers to reduce the IKE processing load on a single key server. Note that only the primary key server sends rekey messages.

When COOP key servers boot, all key servers assume a *secondary* role and begin an election process. One key server, typically the one having the highest priority, is elected as a *primary* key server. The other key servers remain in the secondary state. The primary key server is responsible for creating and distributing group policies to all group members and to periodically synchronize the COOP key servers.

Cooperative key servers exchange one-way announcement messages (primary to secondary). If a secondary key server does not hear from the primary key server for a certain length of time, the secondary key server tries to contact the primary key server and request updated information. If the primary key server does not respond, or if the secondary key server does not hear from the primary key server, a COOP key server reelection is triggered and a new primary key server is elected.

Up to eight key servers can be defined as COOP key servers, but more than four COOP key servers are seldom required. Because rekey information is generated and distributed from a single primary key server, the advantage of deploying more than two key servers is the ability to handle registration load in case of a network failure and reregistration taking place at the same time. This is especially important when using Public Key Infrastructure (PKI) group member authorization because IKE negotiation using PKI requires a lot more CPU power compared to IKE negotiation using pre-shared keys (PSKs).

Tips

- The RSA key must be the same on all cooperative key servers. For information on synchronizing the RSA key, see Generating and Synchronizing RSA Keys, page 29-13.
- It is a best practice to enable periodic ISAKMP keepalives between key servers so that the primary key server can track and display the state of the other secondary key servers. IKE Keepalives between group members and the key server is not required and is not supported. For information on configuring keepalives, see Configuring Global Settings for GET VPN, page 29-16.
- The COOP protocol is configured on a per GDOI group basis. A key server that is configured with multiple GDOI groups can maintain multiple unique COOP relationships with disparate key servers.

Configuring Fail-Close to Protect Registration Failures

Group members must register with the key server to become members of the GET VPN. Before a group member successfully registers with the key server, traffic passing through the group member's GET VPN interface is not encrypted. The period of time in which clear-text transmissions occur can be short (if registration succeeds) or potentially long, if the group member fails to register for any reason.

This default behavior is known as fail-open. If you consider it a violation of your security standards that traffic is sent in clear text at any time, you can configure fail-close mode to protect traffic before (or during) registration. With fail-close mode, all traffic on the interface is dropped except for the traffic you specifically identify in the fail-close ACL. Fail-close mode essentially shuts down the interface until the group member successfully registers with the key server and downloads the required keys and security policy and associations. Note that the use of fail-close mode requires as a minimum Cisco IOS Software release 12.4(22)T or 15.0; you can also configure it on all supported ASRs (see Understanding Devices Supported by Each IPsec Technology, page 25-9).

Fail-close mode is used only during the initial registration. If a group member has already successfully registered, the group member keeps the downloaded policy from the key server even if future registrations fail. However, if you use the **clear crypto gdoi** command on the group member, the subsequent registration attempt is considered a first-time attempt and fail-close mode is enforced.

You configure fail-close mode on the individual group members as described in Configuring GET VPN Group Members, page 29-20. Thus, you can enable the mode on selected group members rather than on all of them. You must specify a fail-close ACL to ensure that you do not lock yourself (and Security Manager) out of the device, preventing configuration updates and maintenance until registration succeeds.

The fail-close ACL is an extended ACL policy object and is configured as part of a crypto map on the device. You configure the rules from the perspective of the group member. Use the following tips to help you create an appropriate fail-close ACL:

- You can configure both **permit** and **deny** statements. In the fail-close ACL, "permit" means "do not send this traffic," whereas "deny" means "send this traffic in clear text." This behavior is different from that of the typical crypto map ACL, where the statements have the following meaning:
  - **Permit**—Means "encrypt this traffic." Because the group member does not have the IPsec security association required to encrypt the traffic prior to registration, the result is that the traffic is dropped.
  - **Deny**—Means "do not encrypt this traffic." In a typical crypto map ACL, a deny statement results in the matching packet being compared to the next crypto map ACL configured on the device (if any). However, if traffic matches a deny statement in the fail-close ACL, all crypto map ACL processing ends and the traffic is allowed in clear text.

The reason deny works this way in fail-close mode is because fail-close includes an implicit ACL statement that gets added at the bottom of the list of crypto map ACLs. This statement is **permit ip any any**, which matches all traffic. Because there is no IPsec security association due to the fact that registration has yet to occur, there is no way to encrypt the remaining traffic and it is dropped.

Note that because of this final permit ip any any statement, you might be able to limit yourself to deny statements in your fail-close ACL.

- The fail-close ACL is processed sequentially after the optional group member security policy ACL. However, all statements in the group member security policy ACL must be deny statements, which indicate that matching traffic should be sent in clear text. Because the security policy is processed according to normal crypto map rules, traffic that matches deny statements is subsequently compared to the fail-close ACL. If the fail-close ACL does not have matching deny statements, the traffic will subsequently be dropped by the implicit final fail-close permit ip any any statement.

  Therefore, if you use a group member security policy ACL, and you want the identified traffic to be sent in clear text regardless of the registration status of the group member, your fail-close ACL should contain all of the same statements contained in the security policy ACL at the least. It might even be possible to use the same ACL object for both ACLs.

  For more information about group member security policies, see Understanding the GET VPN Security Policy and Security Associations, page 29-10.

- The fail-close ACL is inserted as the final crypto map ACL. Thus, if you configure other features on the GET VPN interface that use crypto maps, any traffic identified on deny statements in those other ACLs will also get trapped (and dropped) by the fail-close ACL and the implicit final permit ip any any statement. Thus, configuring fail-close mode for GET VPN can influence the non-GET VPN services you configure on the interface.
- Upon successful registration, the fail-close ACL and the implicit final permit ip any any statement are removed from the crypto maps. These policies are not persistent.
- You should consider including the following rules in the fail-close ACL policy object. Remember that these rules are from the perspective of the group member:
  - SSH, SSL (HTTPS) traffic—You, and Security Manager, need to be able to access the device to configure it. To ensure that you do not lock down the device, include deny statements for SSH and SSL. For SSH, **deny tcp any eq 22 <host or network address>**. For SSL, **deny tcp any eq 443 <host or network address>**. If you specify host addresses, ensure that the Security Manager server is one of the hosts.
  - Routing traffic—To enable routing, allow the traffic for your routing process. For example, if you are using OSPF, **deny ospf any any**.
  - GDOI traffic—Regardless of the contents of the fail-close ACL, the device looks for GDOI registration messages, so you do not need to explicitly allow them to enable successful registration. However, if a group member (1) is in the path between the key server and another group member (2), a registration failure by group member (1) will prevent successful registration by the blocked group member (2). For registration on group member (2) to succeed, the fail-close ACL on group member (1) would have to allow GDOI traffic to pass. Thus, you might want to make it a general practice to allow GDOI traffic in the fail-close ACL: **deny udp any eq 848 any eq 848**.

Related Topics

- Configuring GET VPN, page 29-12
- Creating Access Control List Objects, page 6-53
- Creating Extended Access Control List Objects, page 6-54

Understanding the GET VPN Security Policy and Security Associations

GET VPN uses crypto map access control lists (ACLs) to identify the traffic that needs to be encrypted in the VPN. These ACLs also identify traffic that should be sent as clear text instead of being encrypted (essentially, traffic that lies outside of the VPN). The collection of these ACLs define the security policy for the VPN.

GET VPN provides a multi-layered security policy. You define the general policy for the entire VPN on the key server, but you can also define a separate security policy on group members to account for local variations. The group member security policy always takes priority over the policy received from the key server. When the group member registers with the key server, the group member downloads the key server's security policy and associations and the group member creates a new, single security policy crypto map ACL by concatenating the individual security policies in this order: first, the group member's ACL; second, the key server's first ACL; third, and so forth, any additional ACLs from the key server in the order defined on the key server. It is important to understand that these merged ACLs are treated as a single ACL; they are not searched as separate ACLs. Thus, if traffic matches a deny statement from the group member's ACL, that traffic is never tested against any ACL rules downloaded from the key server.

Tip If a group member leaves the GET VPN, the ACLs downloaded from the key server are removed, but the group member security policy ACL is retained and remains configured on the device.

In GET VPN security policy ACLs (and crypto map ACLs in general), the permit and deny keywords have special meaning:

- **Permit**—Means "encrypt this traffic." Permit entries are allowed only in the security policy ACLs defined on the key server (in the **Group Encryption Policy**), because encrypted traffic needs to have a full IPsec security association, which includes the transform set used for encrypting traffic, and anti-replay and IPsec lifetime configurations. If a packet matches a permit entry, but no IPsec SA exists for that packet, the packet is dropped.

  Normally, your permit rules should be symmetric, that is, the source and destination addresses should be the same. If you need to specify different source and destination addresses, you must create two rules; the second rule should be a mirror image of the first rule, with the source and destination address switched.

- **Deny**—Means "do not encrypt this traffic." In practice, this typically means that the traffic that matches the deny statement is sent in the clear. However, if you configure other features that use crypto maps, "denied" traffic is actually compared to subsequent (lower priority) crypto map ACLs to see if there is a match. IPsec security associations (SAs) are not generated for deny rules.

Following is a summary of the security policies that you can configure, in priority order:

- **Group member security policy**—When you configure the group member, as described in Configuring GET VPN Group Members, page 29-20, you can optionally select an ACL policy object that defines the local group member security policy.

  This group member ACL policy object is allowed to have deny statements only. You use this ACL to identify any traffic that you want to exclude from encryption and send in the clear. For example, if a handful of group members in the group are running a different routing protocol than the usual one, you can configure a local entry to these group members' security policy ACL to bypass encryption of the routing protocol traffic instead of defining the policy globally at the key server level.

- **Key server security policies and security associations**—When you configure the Group Encryption Policy for the GET VPN, as described in Defining GET VPN Group Encryption, page 25-54, you configure ACLs that identify the traffic that should be encrypted and protected in the VPN.

  The security policies on the key server are coupled with transform sets and other settings to define security associations; two IPsec security associations (SAs) are actually configured for every rule within the ACL, and these SAs define how the selected traffic should be encrypted. Thus, all group members use the same group SAs and they do not need to negotiate them with each other.

  Because the key server policy is appended to the group member policy, the policy might be as simple as **permit ip any any**, that is, encrypt all traffic that has not been excluded by the group member policy.

  However, you can create more complex sets of security policies and associations, setting up several separate ACL policy objects that are coupled to different transform sets to define different types of encryption.

  If you create more than one security association, you must identify their order, and they are appended to the group policy in that order. Remember, the end result is a single ACL, so if you include a deny statement in the first ACL, any permit rules for the same traffic in subsequent security associations are ignored, and the traffic is sent in clear text rather than being encrypted.

Note When you consider the security associations defined in the Group Encryption Policy as a whole, you can define up to 100 ACL permit entries. Each permit entry results in a pair of IPsec SAs; the maximum number of IPsec SAs in a group can not exceed 200. It is a best practice to summarize interesting traffic to as few permit entries as possible, and to build symmetric policies, where the source and destination addresses are the same. Unlike traditional IPSec policies, where source and destination address ranges must be uniquely defined, GET VPN is optimized when the source and destination address range are the same. If you configure a rule that has different source and destination addresses, you must also configure the mirrored rule (where the source and destination address are flipped), meaning that four SAs are consumed.

In addition to these security policies, there is an additional fail-close ACL that influences traffic patterns if you configure fail-close mode on a group member. For a complete discussion, see Configuring Fail-Close to Protect Registration Failures, page 29-8.

Related Topics

- Configuring GET VPN, page 29-12
- Creating Access Control List Objects, page 6-53
- Creating Extended Access Control List Objects, page 6-54

Understanding Time-Based Anti-Replay

Anti-replay is an important feature in a data encryption protocol such as IPSec (RFC 2401). Anti-replay prevents a third party from eavesdropping on an IPsec conversation, stealing packets, and injecting those packets into a session at a later time. The time-based anti-replay mechanism helps ensure that invalid packets are discarded by detecting the replayed packets that have already arrived at an earlier time.

GET VPN uses the Synchronous Anti-Replay (SAR) mechanism to provide anti-replay protection for multisender traffic. SAR is independent of real-world Network Time Protocol (NTP) clock or sequential-counter mechanisms (which guarantee packets are received and processed in order). A SAR clock advances regularly. The time tracked by this clock is called pseudotime. The pseudotime is maintained on the key server and is sent periodically to the group members within a rekey message as a timestamp field called pseudoTimeStamp. Group members have to be resynchronized to the pseudotime of the key server periodically. The pseudotime of the key server starts ticking from when the first group member registers. Initially, the key server sends the current pseudotime value of the key server and window size to group members during the registration process. New attributes, such as time-based replay-enabled information, window size, and the pseudotime of the key server, is sent under the SA payload (TEK).

The group members use the pseudotime to prevent replay as follows: the pseudoTimeStamp contains the pseudotime value at which a sender created a packet. A receiver compares the pseudotime value of senders with its own pseudotime value to determine whether a packet is a replayed packet. The receiver uses a time-based anti-replay window to accept packets that contain a timestamp value within that window. The window size is configured on the key server and is sent to all group members.

The following figure illustrates an anti-replay window in which the value PTr denotes the local pseudotime of the receiver, and W is the window size.

You configure anti-replay in the security association definitions in the Group Encryption Policy. For more information, see Defining GET VPN Group Encryption, page 25-54 and Add New or Edit Security Association Dialog Box, page 25-58.

Figure 29-3 Anti-Replay Window

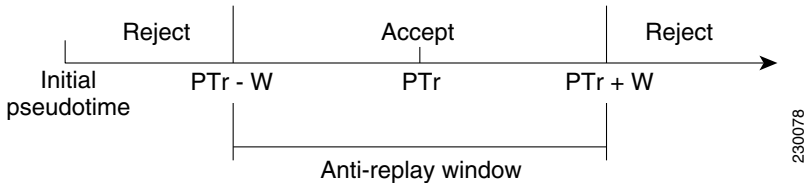

Configuring GET VPN

To configure a full mesh VPN with group encrypted transport (GET), use the Create VPN wizard as described in Creating or Editing VPN Topologies, page 25-28. When you finish the wizard, you are asked if you want to synchronize RSA keys, which is required for normal VPN functioning; for detailed information, see Generating and Synchronizing RSA Keys, page 29-13.

If you select multicast as the rekey transport mechanism, you must enable multicast on all key servers and the desired group members. For more information, see Choosing the Rekey Transport Mechanism, page 29-6.

You can change only the name and description of a GET VPN using the Edit VPN wizard. If you need to make changes to other policies and settings, open the policies from the Site-to-Site Manager page, as follows:

- For ISAKMP and IPSec settings, select **Global Settings for GET VPN**. See Configuring Global Settings for GET VPN, page 29-16.

- For IKE proposal policies, select **IKE Proposal Policy for GET VPN**. See Configuring the IKE Proposal for GET VPN, page 29-15.
- For security associations (ACL rules) and IPSec policies, select **Group Encryption Policy > Security Associations**. See Defining GET VPN Group Encryption, page 25-54.
- For preshared key policies, select **IKEv1 Preshared Key**. See Configuring IKEv1 Preshared Key Policies, page 26-48.
- For public key (PKI) policies, select **Public Key Infrastructure**. See Configuring IKEv1 Public Key Infrastructure Policies in Site-to-Site VPNs, page 26-54.
- For rekey settings, select **Group Encryption Policy > Group Settings**. See Defining GET VPN Group Encryption, page 25-54 and Generating and Synchronizing RSA Keys, page 29-13.
- For key server configuration, including RSA key synchronization, select **Key Servers**. See Configuring GET VPN Key Servers, page 29-18 and Generating and Synchronizing RSA Keys, page 29-13.
- For group membership and endpoint settings, select **Group Members**. See Configuring GET VPN Group Members, page 29-20.

Related Topics

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Understanding the GET VPN Registration Process, page 29-4
- Understanding the GET VPN Security Policy and Security Associations, page 29-10
- Troubleshooting GET VPN Configurations, page 29-25
- Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47

Generating and Synchronizing RSA Keys

When you specify the RSA key label in the Group Encryption Policy (as described in Defining GET VPN Group Encryption, page 25-54), the corresponding RSA key (public and private keys) needs to be configured on all key servers in the GET VPN topology. The key can either be a pre-existing key that you defined on the device, or it could be a new key label, and Security Manager can generate the key for you and synchronize all key servers to use the same key.

You can use the following methods to have Security Manager generate and synchronize the RSA key:

- When creating a new GET VPN using the Create VPN wizard, you are asked at the end of the wizard if you want to synchronize the keys. If you click **Yes**, Security Manager does the key synchronization immediately, and generates a new key if the key does not already exist. For information on using the Create VPN wizard, see Creating or Editing VPN Topologies, page 25-28.
- For an existing GET VPN, you can click the **Synchronize Keys** button on the Key Servers policy. Use this process whenever you add key servers or generate a new key on the primary key server. For information on configuring key server settings for existing topologies, see Configuring GET VPN Key Servers, page 29-18.

Tip For existing GET VPN topologies, if you want to generate a new RSA key, it might be easiest to update the Group Encryption Policy to specify a new, unused RSA key label, then click the Synchronize Keys button in the Key Servers policy. Because the key will not exist on any key server, Security Manager will generate the new key and import it into all key servers. You can then manually delete the old key from each key server.

Following are the uses for the RSA key:

- The key server uses the private RSA key to authenticate rekey messages from the group members.
- The key server provides the public RSA key to group members during registration.
- The key server uses the private key to sign the key encryption key (KEK) and traffic encryption key (TEK). The absence of an RSA key prevents the key server from creating the KEK and TEK.
- The RSA key is also used to sign messages between cooperative key servers.

When you start the RSA key synchronization process, the Synchronize Keys dialog box opens and shows you the overall progress as well as the results for each key server. (You can click the **Abort** button at any time to stop the process.) Security Manager performs the following steps:

1. Logs into all key servers and retrieves the RSA key information from each of them for the RSA key label configured for the VPN.
2. Determines whether any key server has a key with the required label:
   - If no key server has an RSA key with the required label, Security Manager generates the key on the primary key server (the one with the highest priority).
   - If one or more key server does not have the key, but all of the key servers that do have the key have the identical keys, Security Manager uses the existing key on any key server that has it.
   - If more than one key server has the key, but the contents of the key is different among the servers, you are asked if Security Manager can overwrite the keys. If you click **Yes**, Security Manager uses the existing key on the primary key server.

     If you click **No**, you can log into the key servers outside of Security Manager and manually adjust the keys according to your requirements. However, all key servers must have the same key contents for the RSA key. See below for an explanation of the process.
3. Creates an exportable version of the key.
4. Imports the key into each of the remaining key servers.

Tip For the synchronization process to succeed, the devices must be online and reachable and you must have Deploy authorization. If the device connection fails or times out, ensure that you can ping the key server from the Security Manager server. If it is your practice to deploy to file instead of to live devices, you might need to manually generate and synchronize the keys as described below. If you do not have sufficient authorization, you are prevented from initiating the process; someone else must do it.

Manually Generating and Synchronizing the RSA Key

If you do not want Security Manager to generate and synchronize keys, or if for some reason Security Manager cannot complete the process, you can manually generate and synchronize keys using the following sequence in Privileged EXEC (enable) configuration mode:

1. Generate the key on a key server using the following command, where **rekeyrsa** is the name of the key (you can specify a name of your choosing). You must make the key exportable.

crypto key generate rsa general-keys label rekeyrsa modulus 1024 exportable

2. Create an exportable copy of the key using the following command, where **passphrase** is a string used to encrypt the key for import (you can specify your own pass phrase):

   crypto key export rsa rekeyrsa pem terminal 3des passphrase

   This command prints out the public and private keys to the terminal, where you can copy them to the clipboard for import into the other key servers. The keys are demarcated by **----BEGIN/END PUBLIC KEY----** and **----BEGIN/END RSA PRIVATE KEY----**.Note that you can also export to a URL; see the *Cisco IOS Security Command Reference* on Cisco.com for detailed usage information.

3. Import the key into each of the other key servers using the following command:

   crypto key import rsa rekeyrsa pem exportable terminal passphrase

   When copying and pasting the keys, include the begin/end lines.

Configuring the IKE Proposal for GET VPN

Use the IKE Proposal for GET VPN page to define the IKE proposal to be used by the GET VPN topology. The IKE proposal is configured on the key servers and the group members.

These settings are for the ISAKMP security association (SA). If you are using a single key server, the ISAKMP SA is not used after initial group member registration. If you are using more than one key server (cooperative key servers), the ISAKMP SA is needed for communications among the key servers.

To open the IKE Proposal for GET VPN page:

- (Site-to-Site VPN Manager Window) Select an existing GET VPN topology and then select **IKE Proposal for GET VPN** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > IKE Proposal for GET VPN**, and then select an existing policy or create a new one.

The following table explains the settings you can configure in this policy.

Table 29-1 IKE Proposal for GET VPN Policy

Element	Description
IKE Proposal	The IKE proposal policy object that defines the settings you want to use. There are several predefined objects that you might be able to use as is. Click **Select** to open the list of existing IKE proposal objects. The object you select needs to use the same authorization method you are configuring for the group (for example, an object name with the prefix **preshared** when using preshared keys, or with the prefix **cert** when using Public Key Infrastructure (PKI) certificates). When you select an object and click **OK**, the settings defined in the object are displayed in the IKE Proposal Settings display fields. You can also see the settings by editing them in the selection list. If you do not find an appropriate pre-existing object, click the **Add (+)** button in the selection list and create a new object (see Configuring IKEv1 Proposal Policy Objects, page 26-10 for more information and detailed descriptions of the options).

Table 29-1 IKE Proposal for GET VPN Policy (continued)

Element	Description
IKE Proposal Overrides	The number of seconds that the ISAKMP SA for key servers and group members is valid. When the lifetime is exceeded, the SA expires and must be renegotiated between the peers. Values can be 1 to 86400. • If you are using cooperative key servers (more than one key server), set the key server lifetime high. The default 86400 is appropriate. • If you are using a single key server, you can set the lifetime low (but not less than 60 seconds) so that the ISAKMP SA is not retained unnecessarily. It is not used after a group member registers. • We recommend that you set the group member lifetime low as compared to the key server lifetime, especially when cooperative key servers are configured.

Related Topics

- Understanding IKE, page 26-5
- Understanding IKEv1 Preshared Key Policies in Site-to-Site VPNs, page 26-47
- Defining GET VPN Group Encryption, page 25-54
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Configuring Global Settings for GET VPN

Use the Global Settings for GET VPN page to define global settings for ISAKMP and IPsec that apply to devices in your GET VPN topology.

Note The lifetime settings in this policy do not apply to the ISAKMP security association lifetime for the key server and group members. Those lifetime values are configured in the IKE Proposal for GET VPN policy. For more information, see Configuring the IKE Proposal for GET VPN, page 29-15.

To open the Global Settings for GET VPN page:

- (Site-to-Site VPN Manager Window) Select an existing GET VPN topology and then select **Global Settings for GET VPN** in the Policies selector.
- (Policy view) Select **Site-to-Site VPN > Global Settings for GET VPN**, and then select an existing policy or create a new one.

The following table explains the settings you can configure in this policy.

Table 29-2 Global Settings for GET VPN

Element	Description
Enable Keepalive (Key Servers Only)	Whether to enable dead peer detection (DPD) keepalive messages between key servers. If there is more than one key server (cooperative key servers), you should enable periodic keepalive so the servers know each other's status and can elect a new primary server when necessary. Configure the following settings: • **Interval**—When you also select **Periodic**, the number of seconds between DPD messages. If you do not select Periodic, it is the number of seconds during which traffic is not received from the peer before DPD retry messages are sent. The range is from 10 to 3600 seconds. • **Retry**—The number of seconds between DPD retry messages if the DPD retry message is missed by the peer; the range is from 2 to 60 seconds. The default DPD retry message is sent every 2 seconds. Five aggressive DPD retry messages can be missed before the key server is marked as down. • **Periodic**—Whether to send DPD messages at regular intervals (regardless of traffic received from the other key servers). For GET VPN, you should select Periodic.
Identity	During Phase I IKE negotiations, peers must identify themselves to each other. Select the ISAKMP identity to use: • Address—(Default) The IP address of the interface that participates in IKE negotiations. Use the address if only one interface participates in negotiations, and its IP address is known (static). • Hostname—The fully-qualified host name (for example, router1.example.com). • Distinguished Name
SA Requests System Limit	The maximum number of SA requests allowed before IKE starts rejecting them. The specified value must equal or exceed the number of peers, or the VPN tunnels might be disconnected. You can enter a value in the range of 0-99999.
SA Requests System Threshold	The percentage of system resources that can be used before IKE starts rejecting new SA requests. The default is 75 percent.

Table 29-2 Global Settings for GET VPN (continued)

Element	Description
IPsec Settings	Select **Enable Lifetime** if you want to change the default lifetime settings for IPsec SAs. You can configure a lifetime based on the volume of traffic (in kilobytes) between group members, seconds, or both. The key expires when either of the values is reached. The defaults (which are configured even if you do not select this option) are: • **Lifetime (secs)**—3600 seconds (one hour). • **Lifetime (kbytes)**—4,608,000 kilobytes. **Tip** You can override these values for the traffic encryption key when configuring a security association. See Defining GET VPN Group Encryption, page 25-54 and Add New or Edit Security Association Dialog Box, page 25-58.

Related Topics

- Understanding IKE, page 26-5
- Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Configuring GET VPN Key Servers

Use the Key Servers policy to define key servers to be used by a GET VPN topology.

To open the Key Servers policy, in the Site-to-Site VPN Manager Window, select an existing GET VPN topology, then select **Key Servers** from the Policies list.

The table lists the key servers used in the VPN, showing the device name, identity, priority, and registration interface. For detailed information about these attributes, see Edit Key Server Dialog Box, page 29-19.

- To add a key server to the table, click the **Add Row** button and select the device from the list presented. Only devices that can be included as key servers are shown.
- To edit the characteristics of a key server, select it and click the **Edit Row** button. Fill in the Edit Key Server dialog box (see Edit Key Server Dialog Box, page 29-19).
- To delete a key server, select it and click the **Delete Row** button.
- To synchronize the RSA keys among the key servers, so that they all use the identical key, click the **Synchronize Keys** button. For detailed information about the key synchronization process, including when and why you would do it, see Generating and Synchronizing RSA Keys, page 29-13.
- To change the order of a key server when using cooperative key servers, select it and click the up or down arrow button. This order does not define which server is the primary key server (this is determined by the Priority value, the higher the value, the higher the likelihood that the server will be elected the primary key server).

  Instead, the order determines the default order in which group members will try to register with a key server. Group members register with the first key server in the list. If the first key server cannot be reached, group members register with the second key server, and so on. For more information

about key server redundancy, see Configuring Redundancy Using Cooperative Key Servers, page 29-7. Note that you can override this order for individual group members; see Configuring GET VPN Group Members, page 29-20 and Edit Group Member Dialog Box, page 29-21.

Tip You can toggle between showing the interface roles or the actual interfaces defined by those roles in the Identity and interfaces columns using the **Show** field below the table.

Related Topics

- Understanding the GET VPN Registration Process, page 29-4
- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12
- Configuring VPN Topologies in Device View, page 25-19
- Filtering Tables, page 1-48

Add Key Server, Group Member Dialog Box

Use the Add Key Server and Add Group Member dialog boxes to select key servers or group members to be used in the GET VPN topology. Select the check box next to the desired devices and click **OK**.

Navigation Path

To add key servers or group members to a GET VPN topology, click the **Add Row (+)** button beneath the Key Server or Group Member table in the **GET VPN Peers** page of the Create VPN wizard, or for existing topologies, the **Key Servers** or **Group Members** policies. For detailed information, see the following topics:

- Defining GET VPN Peers, page 25-60
- Configuring GET VPN Key Servers, page 29-18
- Configuring GET VPN Group Members, page 29-20

Edit Key Server Dialog Box

Use the Edit Key Servers dialog box to change the attributes defined for a key server in a GET VPN topology.

Navigation Path

- (Create VPN Wizard) Go to the GET VPN Peers Page, select a key server and click the **Edit Row** button. See Defining GET VPN Peers, page 25-60.
- (Site-to-Site VPN Manager Window, page 25-18) Select the **Key Servers** policy, select a key server and click the **Edit Row** button. See Configuring GET VPN Key Servers, page 29-18.

Related Topics

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Field Reference

Table 29-3 Edit Key Server Dialog Box

Element	Description
Identity Interface	The interface that group members use to identify the key server and register with it. The default is the Loopback interface role, which identifies all Loopback interfaces.
Priority	A number between 1-100 that designates the role of the key server, either primary or secondary. The key server with the highest number becomes the primary key server. If two or more key servers are assigned the same priority, the device with the highest IP address is used. The default priority is 100 for the first key server, 95 for the second, and so on. **Note** There can be more than one primary key server if the network is partitioned.
Registration Interface	The interface on which group domain of interpretation (GDOI) registrations can be accepted. If you do not specify a registration interface, GDOI registrations can occur on any interface.

Configuring GET VPN Group Members

Use the Group Members policy to define the group members in a GET VPN topology.

To open the Group Members policy, in the Site-to-Site VPN Manager Window, select an existing GET VPN topology, then select **Group Members** from the Policies list.

The group members table lists the members of the GET VPN, showing the device name, GET-enabled interface, local interface, and security policy. For detailed information about these attributes, see Edit Group Member Dialog Box, page 29-21.

- To add a group member to the table, click the **Add Row** button and select the device from the list presented. Only devices that can be included as group members are shown.
- To edit the endpoint characteristics of a group member, select it and click the **Edit Row** button. Fill in the Edit Group Member dialog box (see Edit Group Member Dialog Box, page 29-21).

  If you select multiple group members in the table, you can also right-click and select the following commands to edit just these attributes:

  - **Edit Key Server Order**—To change the key server list and priority order for the selected group members.
  - **Edit Passive SA Mode**—To change whether the selected group members use passive SA mode.
- To delete a group member, select it and click the **Delete Row** button.

Tip You can toggle between showing the interface roles or the actual interfaces defined by those roles in the interfaces columns using the **Show** field below the table.

Related Topics

- Configuring Fail-Close to Protect Registration Failures, page 29-8
- Using Passive Mode to Migrate to GET VPN, page 29-23

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12
- Configuring VPN Topologies in Device View, page 25-19
- Filtering Tables, page 1-48

Edit Group Member Dialog Box

Use the Edit Group Members dialog box to change the attributes defined for a group member of a GET VPN topology.

Tip If you selected multiple devices and chose an edit command from the right-click menu, this dialog box shows only those options related to the edit command you chose.

Navigation Path

- (Create VPN Wizard) Go to the GET VPN Peers page, select a group member and click the **Edit Row** button. See Defining GET VPN Peers, page 25-60.
- (Site-to-Site VPN Manager Window, page 25-18) Select a GET VPN topology, then select the **Group Members** policy. Select a group member and click the **Edit Row** button. See Configuring GET VPN Group Members, page 29-20.

Related Topics

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Field Reference

Table 29-4 Edit Group Member Dialog Box

Element	Description
GET-Enabled Interface	The VPN-enabled outside interface to the provider edge (PE). Traffic originating or terminating on this interface is evaluated for encryption or decryption, as appropriate. You can configure multiple interfaces. Enter the name of the interface or interface role, or click **Select** to select it from a list or to create a new interface role.
Interface to be used as local address	The interface whose IP address is used to identify the group member to the key server for sending data, such as rekey information. If GET is enabled on only one interface, you do not need to specify the interface to be used as the local address. If GET is enabled on more than one interface, you must specify the interface to be used as the local address. Enter the name of the interface or interface role, or click **Select** to select it from a list or to create a new interface role.

Table 29-4 Edit Group Member Dialog Box (continued)

Element	Description
Security Policy	The local group member security ACL used to deny some group member-specific traffic over and above the security ACL downloaded from the key server. Denied traffic is sent in clear text rather than encrypted. For detailed information, see Understanding the GET VPN Security Policy and Security Associations, page 29-10. Enter the name of the ACL object or click **Select** to select it from a list or to create a new object.
Enable Fail Close Fail Close ACL	Whether to enable fail-close mode on the device, which prevents the device from transmitting clear text traffic before the device successfully registers with the key server. Fail-close mode requires as a minimum Cisco IOS Software release 12.4(22)T or 15.0; you can also configure it on all supported ASRs. **Tip** Fail-close mode is a complex feature, and you must carefully construct the fail-close ACL or you might lock yourself out of the device. Before enabling fail-close mode, read Configuring Fail-Close to Protect Registration Failures, page 29-8. You must select an ACL policy object that identifies allowable clear text traffic (using deny statements), such as SSH and SSL communications with the Security Manager server to allow for configuration updates. Enter then name of the object or click **Select** to select it or to create a new object.
Override Key Servers	Whether to override the key server list configured for the GET VPN topology as a whole for this particular group member. If you select this option, you can choose a subset of the key servers configured for the topology to be used by the selected group member, and change their priority order. This can help you load-balance registration activity among a group of cooperative key servers. For more information, see Configuring Redundancy Using Cooperative Key Servers, page 29-7. Click **Select** to change the key server list and priority order of the key servers using the Key Servers Selection dialog box. A key server must be defined for the GET VPN topology before you can modify its use for a group member.

Table 29-4 Edit Group Member Dialog Box (continued)

Element	Description
Enable Passive SA Mode	Whether to put the group member into passive security association (SA) mode, which means the group member installs the SA in the inbound direction only. This means the group member can receive encrypted data, but it sends clear text data only. This mode is useful for testing the VPN only, primarily when you are migrating from an existing VPN to a GET VPN. (The group member must be running Cisco IOS Software version 12.4(22)T or 15.0 at minimum, or be a supported ASR, to use this mode.) This setting is similar to the **Receive Only** setting in the Group Encryption Policy, which applies to the topology as a whole. This group member option overrides the setting in the Group Encryption Policy. For detailed information on how you can use these passive mode features to migrate or test a GET VPN, see Using Passive Mode to Migrate to GET VPN, page 29-23.

Using Passive Mode to Migrate to GET VPN

If you are migrating an existing VPN to the GET VPN technology, especially a clear-text VPN, you can use two features to help you migrate in a phased approach to help prevent network down-time. The features are essentially the same, and involve the passive acceptance of encrypted traffic, but you configure them on different devices in the GET VPN.

Normally, in a fully-deployed GET VPN, traffic is encrypted in both directions (bidirectional security associations, or SAs). However, during testing, you can use passive mode. In passive mode, the group member installs the SA in the inbound direction only, so that the group member receives encrypted traffic but sends traffic in clear text. You can then test the VPN to ensure that it is performing as expected before turning on full encryption.

Use the following features to configure passive mode in a GET VPN:

- **SA Receive-only mode**—You configure receive-only mode for security associations on the key servers in the topology using the Group Encryption Policy. Thus, the setting applies to the entire topology.
- **Passive SA mode**—You configure passive security association mode on individual group members. This setting overrides the SA receive-only setting; thus, you can turn on full encryption for the entire topology, but leave some group members in passive mode. This lets you test the group members in stages and enable full encryption after you verify each member device.

Tip Passive SA mode on group members requires Cisco IOS Software release 12.4(22)T+ or 15.0+, or Release 2.3 (12.2(33)XNC)+ on ASRs.

The following procedure shows an example of the end-to-end migration process you might follow to convert to GET VPN using these passive mode features.

Related Topics

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

Step 1 Create the new GET VPN topology in Security Manager using the Create VPN wizard. When you are in the wizard, ensure that you make these selections:

- When selecting devices, choose the key servers for the topology, but for group members, select the first set of group members that will be migrated. For more information, see Selecting Devices for Your VPN Topology, page 25-32.
- When configuring the group encryption settings, select **Receive Only**. This enables the SA receive-only feature for the entire topology. For more information, see Defining GET VPN Group Encryption, page 25-54.

For information about creating VPNs, see Creating or Editing VPN Topologies, page 25-28.

Step 2 Deploy the configurations to all devices in the VPN. The group members should now be able to receive encrypted traffic but not send it. For information on the deployment process, see the following topics based on the Workflow mode you are using:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Step 3 Outside of Security Manager, verify that all of the group members are functioning properly.

For example, you can test whether the group members are able to send and receive encrypted packets using some CLI commands on the group member devices:

- On group member 1, configure the following command, where “groupexample” is the name of the GDOI group for the VPN. This command sets the device to accept encrypted or clear text, but to send only clear text.

  crypto gdoi gm group groupexample ipsec direction inbound only

- On group member 2, configure the following command. This command sets the device to accept encrypted or clear text, but to send encrypted text.

  crypto gdoi gm group groupexample ipsec direction inbound optional

- Ping group member 1 from group member 2. Group member 2 should encrypt the packet before sending it, and group member 1 should accept it and decrypt it. If you ping member 2 from member 1, the ping should be sent in clear text and accepted by member 2. Ensure that your ACLs allow pings.

Step 4 In Security Manager, select **Manage > Site-to-Site VPNs** (see Site-to-Site VPN Manager Window, page 25-18).

Select the GET VPN topology, then select **Group Members**.

Add the remaining group members that you want to add to the topology (click the **Add Group Member (+)** button, select the devices, and click **OK**).

If you want to use passive mode to test the new group members before enabling full encryption, ensure that you select **Enable Passive SA Mode** when configuring the group members:

- To configure an individual group member, select it and click the **Edit Group Member (pencil)** button.
- To enable passive mode on more than one device at a time, use Shift+click or Ctrl+click to select multiple devices, then right-click and select **Edit Passive SA Mode**. You can then select the option and click **OK**.

For more information on configuring group members, see Configuring GET VPN Group Members, page 29-20.

Step 5 Deploy the configuration changes to all devices in the VPN. All devices should be operating in passive mode at this point.

Step 6 In the Site-to-Site VPN Manager, select the GET VPN topology, then select **Group Encryption Policy**.

Deselect **Receive Only**. This turns off SA receive-only mode at the topology level.

Step 7 Deploy the configuration changes to all devices in the VPN. Now the GET VPN should be operating in fully encrypted mode for the original group members that you tested. Any new members that you added with passive SA mode enabled should be receiving encrypted traffic and sending clear text traffic.

Step 8 Use the following process to verify the new devices and to turn off passive mode. You can follow this process for all new devices at once, or you can do smaller groups of them at a time. You can also use this process for new group members as you extend your network. Iterate the following steps as appropriate:

a. Verify that the new group members are functioning properly using the same techniques that you used to verify the original group members.

b. When you are ready to move a set of group members to fully-encrypted mode, in the Site-to-Site VPN Manager, select the GET VPN topology and select **Group Members**.

c. Select all passive mode group members that should use full encryption, right-click and select **Edit Passive SA Mode**. Deselect the **Enable Passive SA Mode** option and click **OK**.

d. Deploy configurations to all devices in the VPN, not just the ones whose passive mode you changed. Normally, you should not deploy to less than all devices in a VPN.

Troubleshooting GET VPN Configurations

If after provisioning and deploying GET VPN using Security Manager, the GET VPN is not working, check the following:

- Ensure that the RSA key is synchronized among all cooperative key servers (that is, the RSA key is the same). For information on how to synchronize keys, see Generating and Synchronizing RSA Keys, page 29-13.
- If desired traffic is not being encrypted, make sure the key server security policy ACL (security association) has a permit ACE for the desired traffic. For asymmetric ACEs (where the source and destination addresses are different), ensure that there is a mirrored ACE (with the source and destination addresses reversed). For more information, see Understanding the GET VPN Security Policy and Security Associations, page 29-10.
- For multicast rekey, make sure that the network is multicast enabled and that all key servers and most group members are configured to enable multicast. You must enable multicast on the devices directly; Security Manager does not provision the commands required to enable multicast. For more information, see Choosing the Rekey Transport Mechanism, page 29-6.
- When using multicast rekey, check whether there is a deny ACE in the key server security ACL for the multicast group address to prevent encryption of multicast rekey messages.
- Check that the local security ACL on the group member has only deny ACEs. If you include a permit statement in an attempt to identify traffic that should be encrypted, the matching traffic is actually dropped because there is no corresponding IPsec SA. Because the permit entry is defined in the group member, the key server is not aware of it and cannot generate the required IPsec SA. For more information, see Understanding the GET VPN Security Policy and Security Associations, page 29-10.
- For group member authorization using certificates, check that the ISAKMP authentication uses certificates and that a PKI policy is configured. ISAKMP identity on the group member and key server should be set to use the distinguished name (dn).

- Normally, network address translation (NAT) is not used in the type of WAN environments where GET VPN is deployed. However, if you use NAT, ensure that the security policy ACL has permit statements for the translated addresses. Also, if you are using Network Address Translation-Traversal (NAT-T), the GDOI protocol port changes to 4500.
- A control plane replay protection mechanism was added to Cisco IOS Software releases 12.4(15)T10, 12.4(22)T3, 12.4(24)T2, 15.0(1)M, and 12.2(33)XNE. This mechanism is not backward-compatible, so if any GET VPN group member in the network is running any of these (or later) releases, you must also upgrade all key servers to one of these (or newer) releases. Otherwise, network disruption might occur because of a failed rekey, which causes one of the following system logging (syslog) messages to appear:
  - %GDOI-3-GDOI_REKEY_SEQ_FAILURE: Failed to process rekey seq # 2 in seq payload for group get-group, last seq # 6
  - %GDOI-3-PSEUDO_TIME_TOO_OLD: Rekey received in group get-group is too old and failed PST check: my_pst is 184 sec, peer_pst is 25 sec, allowable_skew is 10 sec

Tip For additional troubleshooting tips from the CLI configuration perspective, including information about valuable **show** commands, see *Cisco Group Encrypted Transport VPN* on Cisco.com.

Related Topics

- Understanding Group Encrypted Transport (GET) VPNs, page 29-2
- Configuring GET VPN, page 29-12

CHAPTER 30

Managing Remote Access VPNs: The Basics

Cisco Security Manager lets you configure both remote access IPSec VPNs and remote access SSL VPNs. Security Manager provides flexible configuration and management of remote access VPNs:

- You can discover existing remote access VPN configuration policies from existing live devices or from configuration files. Then, you can change and deploy new or updated policies, as necessary.
- You can use the configuration wizard to help you quickly and easily set up these two types of remote access VPNs with basic functionality.
- If you know the functions and feature your network requires, you can configure remote access VPNs independently. You can also use the wizard to create a basic remote access VPN and then configure additional features that are not included in the wizard separately.

In addition, Cisco Security Manager provides flexibility in how remote access VPN configuration policies are assigned: Device view or Policy view.

For some policies, you can also assign either the factory default policy (a private policy), or a shared policy that you created using Security Manager.

This chapter contains the following topics:

- Understanding Remote Access VPNs, page 30-1
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8
- Overview of Remote Access VPN Policies, page 30-9
- Discovering Remote Access VPN Policies, page 30-12
- Using the Remote Access VPN Configuration Wizard, page 30-13

Understanding Remote Access VPNs

Security Manager supports two types of remote access VPNs: IPSec and SSL.

This section contains the following topics:

- Understanding Remote Access IPSec VPNs, page 30-2
- Understanding Remote Access SSL VPNs, page 30-2

Understanding Remote Access IPSec VPNs

Remote access IPSec VPNs permit secure, encrypted connections between a company's private network and remote users, by establishing an encrypted IPSec tunnel across the Internet using broadband cable, DSL, dial-up, or other connections.

A remote access IPSec VPN consists of a VPN client and a VPN headend device, or VPN gateway. The VPN client software resides on a user's workstation and initiates the VPN tunnel access to the corporate network. At the other end of the VPN tunnel is the VPN gateway at the edge of the corporate site.

When a VPN client initiates a connection to the VPN gateway device, negotiation consists of authenticating the device through Internet Key Exchange (IKE), followed by user authentication using IKE Extended Authentication (Xauth). Next the group profile is pushed to the VPN client using mode configuration, and an IPsec security association (SA) is created to complete the VPN connection.

Tip For a remote access IPsec VPN hosted on an ASA 8.4(x) device, you have the option of configuring IKE version 2 (IKEv2). If you decide to use IKEv2, you must configure several SSL VPN policies in addition to the regular IPSec policies. The user also must use the AnyConnect 3.0+ VPN client to make an IKEv2 connection. For more information, see Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25.

For remote access IPSec VPNs, AAA (authentication, authorization, and accounting) is used for secure access. With user authentication, a valid user name and password must be entered before the connection is completed. User names and passwords can be stored on the VPN device itself, or on an external AAA server that can provide authentication to numerous other databases. For more information on using AAA servers, see Understanding AAA Server and Server Group Objects, page 6-27.

Note Site-to-site Easy VPN topologies use some of the same policies and policy objects that are used in remote access IPsec VPNs, but the policies are kept distinct from the remote access policies. In Easy VPN, the remote clients are hardware clients, such as routers, whereas in remote access IPSec VPNs, remote clients are workstations or other devices that use VPN client software. For more information, see Understanding Easy VPN, page 28-1.

Related Topics

- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25
- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices), page 30-36
- Overview of Remote Access VPN Policies, page 30-9
- Discovering Remote Access VPN Policies, page 30-12

Understanding Remote Access SSL VPNs

An SSL VPN lets users access enterprise networks from any Internet-enabled location. Users can make clientless connections, which use only a Web browser that natively supports Secure Socket Layer (SSL) encryption, or they can make connections using a full client (such as AnyConnect) or a thin client.

Note SSL VPN is supported on ASA 5500 devices running software version 8.0 and later, running in single-context and router modes, on Cisco 870, 880, 890, 1800, 2800, 3700, 3800, 7200, and 7301 Series routers running software version 12.4(6)T and later, and on Cisco 1900, 2900, and 3900 Series routers running software version 15.0(1)M and later. For the 880 Series routers, the minimum software version is 12.4(15)XZ, which is mapped to 12.4(20)T in Security Manager.

On IOS devices, remote access is provided through an SSL-enabled VPN gateway. Using an SSL-enabled Web browser, the remote user establishes a connection to the SSL VPN gateway. After the remote user is authenticated to the secure gateway via the Web browser, an SSL VPN session is established and the user can access the internal corporate network. A portal page lets users access all the resources available on the SSL VPN networks.

On ASA devices, remote users establish a secure, remote access VPN tunnel to the security appliance using the Web browser. The SSL protocol provides the secure connection between remote users and specific, supported internal resources that you configure at a central site. The security appliance recognizes connections that need to be proxied, and the HTTP server interacts with the authentication subsystem to authenticate users.

User authentication can be done using usernames and passwords, certificates, or both.

Note Network administrators provide user access to SSL VPN resources on a group basis instead of on an individual user basis.

This section contains the following topics:

- Remote Access SSL VPN Example, page 30-3
- SSL VPN Access Modes, page 30-4
- Understanding and Managing SSL VPN Support Files, page 30-5
- Prerequisites for Configuring SSL VPNs, page 30-7
- SSL VPN Limitations, page 30-8

Remote Access SSL VPN Example

The following illustration shows how a mobile worker can access protected resources from the main office and branch offices. Site-to-site IPsec connectivity between the main and remote sites is unaltered. The mobile worker needs only Internet access and supported software (Web browser and operating system) to securely access the corporate network.

Figure 30-1 Secure SSL VPN Access Example

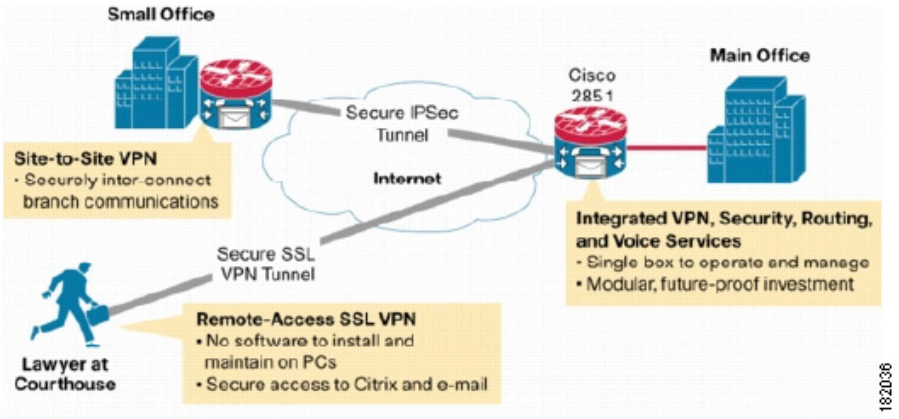

SSL VPN Access Modes

SSL VPN provides three modes of remote access on IOS routers: Clientless, Thin Client and Full Client. On ASA devices, there are two modes: Clientless (which includes Clientless and Thin Client port forwarding) and AnyConnect Client (a full client).

Clientless Access Mode

In Clientless mode, the remote user accesses the internal or corporate network using a Web browser on the client machine. No applet downloading is required.

Clientless mode is useful for accessing most content that you would expect in a Web browser, such as Internet access, databases, and online tools that employ a Web interface. It supports Web browsing (using HTTP and HTTPS), file sharing using Common Internet File System (CIFS), and Outlook Web Access (OWA) email. For Clientless mode to work successfully, the remote user's PC must be running Windows 2000, Windows XP, or Linux operating systems.

Browser-based SSL VPN users connecting from Windows operating systems can browse shared file systems and perform the following operations: view folders, view folder and file properties, create, move, copy, copy from the local host to the remote host, copy from the remote host to the local host, and delete. Internet Explorer indicates when a Web folder is accessible. Accessing this folder launches another window, providing a view of the shared folder, on which users can perform web folder functions, assuming the properties of the folders and documents permit them.

Thin Client Access Mode

Thin Client mode, also called TCP port forwarding, assumes that the client application uses TCP to connect to a well-known server and port. In this mode, the remote user downloads a Java applet by clicking the link provided on the portal page. The Java applet acts as a TCP proxy on the client machine for the services configured on the SSL VPN gateway. The Java applet starts a new SSL connection for every client connection.

The Java applet initiates an HTTP request from the remote user client to the SSL VPN gateway. The name and port number of the internal email server is included in the HTTP request. The SSL VPN gateway creates a TCP connection to that internal email server and port.

Thin Client mode extends the capability of the cryptographic functions of the Web browser to enable remote access to TCP-based applications such as Post Office Protocol version 3 (POP3), Simple Mail Transfer Protocol (SMTP), Internet Message Access protocol (IMAP), Telnet, and Secure Shell (SSH).

Note The TCP port-forwarding proxy works only with Sun's Java Runtime Environment (JRE) version 1.4 or later. A Java applet is loaded through the browser that verifies the JRE version. The Java applet refuses to run if a compatible JRE version is not detected.

When using Thin Client mode, you should be aware of the following:

- The remote user must allow the Java applet to download and install.
- For TCP port-forwarding applications to work seamlessly, administrative privileges must be enabled for remote users.
- You cannot use Thin Client mode for applications such as FTP, where the ports are negotiated dynamically. That is, you can use TCP port forwarding only with static ports.

Full Tunnel Client Access Mode

Full Tunnel Client mode enables access to the corporate network completely over an SSL VPN tunnel, which is used to move data at the network (IP) layer. This mode supports most IP-based applications, such as Microsoft Outlook, Microsoft Exchange, Lotus Notes E-mail, and Telnet. Being part of the SSL VPN is completely transparent to the applications run on the client. A Java applet is downloaded to handle the tunneling between the client host and the SSL VPN gateway. The user can use any application as if the client host was in the internal network.

The tunnel connection is determined by the group policy configuration. The SSL VPN client (SVC) or AnyConnect client is downloaded and installed to the remote client, and the tunnel connection is established when the remote user logs in to the SSL VPN gateway. By default, the client software is removed from the remote client after the connection is closed, but you can keep it installed, if required.

Note Full Tunnel SSL VPN access requires administrative privileges on the remote client.

Understanding and Managing SSL VPN Support Files

SSL VPNs sometimes require supporting files that reside in the device's flash storage. This is especially true of SSL VPNs configured on ASA devices. Supporting files include Cisco Secure Desktop (CSD) packages, AnyConnect client images, and plug-in files. Security Manager includes many of these files for your use. However, some supporting files, such as graphic files used for portal pages, or client profiles used for AnyConnect clients are not provided by Security Manager.

Typically, you need to create a File Object to specify a supporting file, and you then select the File Object when you create a policy that refers to it. You can create the File Objects that you need when you create the policies, or you can create them before you start defining policies. For more information, see Add and Edit File Object Dialog Boxes, page 34-37.

When you deploy policies to the devices, any supporting files referenced in your policies are copied to the device and placed in flash memory in the \csm folder. For the most part, you do not have to do any manual work to make this happen. The following are some situations where you might need to do some manual work:

- If you are trying to discover existing SSL VPN policies, or rediscover them, file references from the SSL VPN policies must be correct. For detailed information on how supporting files are handled during policy discovery, see Discovering Remote Access VPN Policies, page 30-12.
- If you have configured the ASA device in an Active/Failover configuration, you must get the supporting files onto the failover device. The supporting files are not copied over to the failover device during a failover. You have these choices for getting the files onto the failover device:

  - Manually copy the files from the \csm folder on the active unit to the failover unit.
  - After deploying the policies to the active unit, force a failover and redeploy the policies to the now-active unit.
- If you are using a VPN cluster for load balancing, the same supporting files must be deployed to all devices in the cluster.

Cisco Secure Desktop (CSD) Packages

These packages are for ASA SSL VPNs. You select a package in the Dynamic Access policy. The package you select must be compatible with the ASA operating system version running on the device. When you create a Dynamic Access policy for an ASA device, the version number that is compatible with the device's operating system is displayed in the Version field.

You can find the CSD packages in Program Files\CSCOpx\files\vms\repository\. The file names are in the form securedesktop-asa_k9-*version*.pkg or csd_*version*.pkg, where *version* is the CSD version number such as 3.5.1077.

Following is the CSD compatibility with ASA versions for the CSD packages shipped with Security Manager:

- csd_3_6_181-3.6.181.pkg—ASA 8.4 or later.
- csd_3_5_2008-3.5.2008.pkg—ASA 8.0(4) or later.
- csd_3_5_2001-3.5.2001.pkg—ASA 8.0(4) or later.
- csd_3_5_1077-3.5.1077.pkg—ASA 8.0(4) or later.
- csd_3_5_841-3.5.841.pkg—ASA 8.0(4) or later.
- csd_3_4_2048-3.4.2048.pkg—ASA 8.0(4) or later.
- csd_3_4_1108-3.4.1108.pkg—ASA 8.0(4) or later.
- securedesktop_asa_k9-3.3.0.151.pkg—ASA 8.0(3.1) or later.
- securedesktop_asa-k9-3.3.0.118.pkg—ASA 8.0(3.1) or later.
- securedesktop-asa-k9-3.2.1.126.pkg—ASA 8.0(3) or later.
- securedesktop-asa_k9-3.2.0.136.pkg—ASA 8.0(2) or later.

For more information on CSD version compatibility with ASA versions, see the CSD release notes at http://www.cisco.com/en/US/products/ps6742/prod_release_notes_list.html and Supported VPN Platforms on Cisco.com.

For more information on creating Dynamic Access policies to specify the CSD, see Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9.

AnyConnect Client Images

These images are for remote access SSL and IKEv2 IPsec VPNs hosted on an ASA. The AnyConnect client is downloaded to the user's PC and manages the client's VPN connection. Security Manager includes several AnyConnect images, which you can find in Program Files\CSCOpx\files\vms\repository\. The package names indicate the workstation operating system and the anyconnect release number in this general pattern: anyconnect-*client_OS_information-anyconnect_release*.pkg. For example, anyconnect-win-3.0.0610-k9-3.0.0610.pkg is the AnyConnect 3.0(0610) client for Windows workstations. The k9 indicates that the package includes encryption. In this example, the AnyConnect release number is repeated; in some file names, the release number appears once.

Packages are available for the following workstation operating systems (OS). For specific information on which OS versions that each client supports, see the documentation for the AnyConnect client on Cisco.com.

- Linux—Packages start with anyconnect-linux, or anyconnect-linux-64 for 64-bit versions.
- Mac OS—Packages start with anyconnect-macosx for Mac OS X on i386 workstations, and anyconnect-macosx-powerpc for Mac OS X on Power PC workstations.
- Windows—Packages start with anyconnect-win.

You can also download other AnyConnect client packages to the Security Manager server or your local Security Manager client and use them in remote access policies. However, Security Manager might not be able to configure newer parameters for those clients, although it might be possible to use FlexConfigs to configure newer parameters.

For more information on the AnyConnect client, its profiles, and how to configure policies to load the client onto the device, see the following topics:

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Cisco AnyConnect Profile Editor, page 31-63

Plug-in Files

These files are used as browser plug-ins. You can find plug-in files in Program Files\CSCOpx\files\vms\repository\. For complete information on the available files, see Configuring SSL VPN Browser Plug-ins (ASA), page 31-60.

Prerequisites for Configuring SSL VPNs

For a remote user to securely access resources on a private network behind an SSL VPN gateway, the following prerequisites must be met:

- A user account (login name and password).
- An SSL-enabled browser (such as Internet Explorer, Netscape, Mozilla, or Firefox).
- An email client (such as Eudora, Microsoft Outlook, or Netscape Mail).
- One of the following operating systems:
  - Microsoft Windows 2000 or Windows XP, with either JRE for Windows version 1.4 or later, or a browser that supports ActiveX controls.
  - Linux with JRE for Linux version 1.4 or later. To access Microsoft shared files from Linux in clientless remote access mode, Samba must also be installed.

Related Topics

- SSL VPN Access Modes, page 30-4
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32

SSL VPN Limitations

SSL VPN configurations in Security Manager are subject to the following limitations:

- SSL VPN license information cannot be imported into Security Manager. As a result, certain command parameters, such as **vpn sessiondb** and **max-webvpn-session-limit**, cannot be validated.
- You must configure DNS on each device in the topology in order to use clientless SSL VPN. Without DNS, the device cannot retrieve named URLs, but only URLs with IP addresses.
- If you share your Connection Profiles policy among multiple ASA devices, bear in mind that all devices share the same address pool unless you use device-level object overrides to replace the global definition with a unique address pool for each device. Unique address pools are required to avoid overlapping addresses in cases where the devices are not using NAT.
- If the device configuration contains an address pool for SSL VPN with a name that begins CSM_ (the naming convention used by Cisco Security Manager), Cisco Security Manager cannot detect whether the addresses in that pool overlap with the pool configured in your SSL VPN policy. (This can occur, for example, when the pool was configured by a user on a different installation of Security Manager.) This can lead to errors during deployment. Therefore, we recommend that you configure the same IP address pool as a network/host object in Security Manager and define it as part of the SSL VPN policy. This enables the proper validation to take place.
- The same IP address and port number cannot be shared by multiple SSL VPN gateways on the same IOS device. As a result, deployment errors can occur if a duplicate gateway exists in the device configuration but was not redefined using the Security Manager interface. If such an error occurs, you must choose a different IP address and port number and redeploy.
- If you define AAA authentication or accounting as part of an SSL VPN policy, the **aaa new-model** command is deployed to enable AAA services. Bear in mind that this command is not removed if you later delete the SSL VPN policy, as there might be other parts of the device configuration that require the **aaa new-model** command for AAA services.

Note In addition, we recommend that you define at least one local user on the device with a privilege level of 15. This ensures that you will not be locked out of the device if the **aaa new-model** command is configured without an associated AAA server.

Related Topics

- SSL VPN Access Modes, page 30-4
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32

Understanding Devices Supported by Each Remote Access VPN Technology

There are three types of remote access VPN: IKE version 1 (IKEv1) IPsec, IKE version 2 (IKEv2) IPsec, and SSL. The devices on which you can configure these technologies differs, and broadly speaking, the configuration for each type of VPN differs for ASA/PIX 7.0+ compared to IOS/PIX 6.3 devices.

The following table describes the basic device support. When you select a device, the device type will determine which remote access policies are visible or configurable.

Tip Some device models have NO-VPN versions, which do not support VPN configuration. Thus, although the 3845 model might be supported for a type of VPN, the 3845 NOVPN model is not supported. In addition, the Cisco Catalyst 6500 series ASA Services Module (running software release 8.5(x)) does not support any type of VPN.

Table 30-1 Devices Supported by Each Remote Access Technology

Technology	Supported Platforms
IKE version 1 IPsec	• ASA/PIX 7.0+—ASA 5500 series and PIX 515, 515E, 525, or 535 with PIX Software 7.0+ (including 8.0+), running in single context and router modes. • IOS/PIX 6.3—Cisco IOS security routers (including Aggregation Service Routers, or ASRs), Catalyst 6500/7600, and PIX Firewalls running PIX Software 6.3 only.
IKE version 2 IPsec	ASA 5500 series only, running ASA Software 8.4(x) only.
SSL	• ASA—ASA 5500 series devices running software version 8.0 and later, running in single-context and router modes. • IOS—Cisco 870, 880, 890, 1800, 2800, 3700, 3800, 7200, and 7301 Series routers running software version 12.4(6)T and later, and on Cisco 1900, 2900, and 3900 Series routers running software version 15.0(1)M and later. For the 880 Series routers, the minimum software version is 12.4(15)XZ, which is mapped to 12.4(20)T in Security Manager. **Tip** No version of PIX is supported for SSL VPN configuration.

Related Topics

- Understanding Remote Access IPSec VPNs, page 30-2
- Understanding Remote Access SSL VPNs, page 30-2
- Using the Remote Access VPN Configuration Wizard, page 30-13
- Overview of Remote Access VPN Policies for ASA and PIX 7.0+ Devices, page 31-2
- Overview of Remote Access VPN Policies for IOS and PIX 6.3 Devices, page 33-2

Overview of Remote Access VPN Policies

The following list summarizes the various policies used in remote access VPN configuration based on the technology used in the VPN. Possible remote access VPN types are: IKE version 1 (IKEv1) IPsec, IKE version 2 (IKEv2) IPsec, and SSL. Where indicated, many of these policies apply to specific device types only. To see an edited version of this list per device type, see the following topics:

- Overview of Remote Access VPN Policies for ASA and PIX 7.0+ Devices, page 31-2
- Overview of Remote Access VPN Policies for IOS and PIX 6.3 Devices, page 33-2

Note You cannot configure SSL VPNs on PIX devices; PIX devices support remote access IKEv1 IPsec VPNs only.

Note You can create Unified ACL object on-the-fly in certain Remote Access VPN policies, such as Dynamic Access Policy. However, when you create Unified ACL object on-the-fly, Cisco Security Manager displays an error message. To overcome this issue, you must select the created ACL in the Selector window and save the policy.

- **Policies used with remote access IKEv1 and IKEv2 IPsec and SSL VPNs:**
  - **ASA Cluster Load Balancing** (ASA/PIX 7.0+)—In a remote client configuration in which you are using two or more devices connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called load balancing. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. Load balancing is effective only on remote sessions initiated with an ASA device. For more information, see Understanding Cluster Load Balancing (ASA), page 31-5.
  - **Connection Profiles** (ASA/PIX 7.0+)—A connection profile is a set of records that contain VPN tunnel connection policies, including the attributes that pertain to creating the tunnel itself. Connection profiles identify the group policies for a specific connection, which includes user-oriented attributes. For more information, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.
  - **Dynamic Access** (ASA 8.0+)—Multiple variables can affect each VPN connection, for example, intranet configurations that frequently change, the various roles that each user might inhabit within an organization, and logins from remote access sites with different configurations and levels of security. Dynamic access policies (DAP) let you configure authorization that addresses these many variables. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. For more information, see Chapter 32, "Managing Dynamic Access Policies for Remote Access VPNs (ASA 8.0+ Devices)".
  - **Global Settings**—You can define global settings that apply to all devices in your remote access VPNs. These settings include Internet Key Exchange (IKE), IKEv2, IPsec, NAT, and fragmentation definitions. The global settings typically have defaults that work in most situations, so configuring the Global Settings policy is optional in most cases; configure it only if you need non-default behavior or if you are supporting IKEv2 negotiations. For more information, see Configuring VPN Global Settings, page 26-30.
  - **Group Policies** (ASA/PIX 7.0+)—You can view the user group policies defined for your remote access VPN connection profiles. From this page, you can specify new ASA user groups and edit existing ones. When you create a connection profile, if you specify a group policy that has not been used on the device, the group policy is automatically added to the Group Policies page; you do not need to add it to this policy before you create the connection profile. For more information, see Configuring Group Policies for Remote Access VPNs, page 31-26.
  - **Public Key Infrastructure**—You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and to manage keys and certificates. Certification Authority (CA) servers are used to manage these certificate requests and issue certificates to users who connect to your IPsec or SSL remote access VPN. For more information, see Understanding Public Key Infrastructure Policies, page 26-51 and Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.

- **Policies used in remote access IPsec VPNs only:**
  - **Certificate To Connection Profile Maps, Policy and Rules** (IKEv1 IPSec only, ASA/PIX 7.0+ only.)—Certificate to connection profile map policies let you define rules to match a user's certificate to a permission group based on specified fields. To establish authentication, you can use any field of the certificate, or you can have all certificate users share a permission group. You can match the group from the DN rules, the Organization Unit (OU) field, the IKE identity, or the peer IP address. You can use any or all of these methods. For more information, see Configuring Certificate to Connection Profile Map Policies (ASA), page 31-36.
  - **IKE Proposal**—Internet Key Exchange (IKE), also called ISAKMP, is the negotiation protocol that enables two hosts to agree on how to build an IPsec security association. IKE is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs). Use the IKE Proposal policy to define the requirements for phase 1 of the IKE negotiation. For more information, see Configuring an IKE Proposal, page 26-9.
  - **IPsec Proposal (ASA/PIX 7.x)**—An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA. The policy is used for IKE phase 2 negotiations. For more information, see Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40.
  - **IPsec Proposal (IOS/PIX 6.x)**—An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA. The policy is used for IKE phase 2 negotiations. For more information, see Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3.
  - **High Availability** (IOS/PIX 6.3)—High Availability (HA) is supported by the creation of an HA group made up of two or more hub devices that use Hot Standby Routing Protocol (HSRP) to provide transparent, automatic device failover. For more information, see Configuring High Availability in Remote Access VPNs (IOS), page 33-11.
  - **User Groups (IOS/PIX 6.x)**—A user group policy specifies the attributes that determine user access to and use of the VPN. For more information, see Configuring User Group Policies, page 33-13.
- **Policies used in remote access IKEv2 IPSec and SSL VPNs only:**
  - **Access** (ASA only.)—An Access policy specifies the security appliance interfaces on which a remote access SSL or IKEv2 IPsec VPN connection profile can be enabled, the port to be used for the connection profile, Datagram Transport Layer Security (DTLS) settings, the SSL VPN session timeout and maximum number of sessions. You can also specify whether to use the AnyConnect VPN Client or AnyConnect Essentials Client. For more information, see Understanding SSL VPN Access Policies (ASA), page 31-44.
  - **Other Settings** (ASA only.)—The SSL VPN Other Settings policy defines settings that include caching, content rewriting, character encoding, proxy and proxy bypass definitions, browser plug-ins, AnyConnect client images and profiles, Kerberos Constrained Delegation, and some other advanced settings. For more information, see Configuring Other SSL VPN Settings (ASA), page 31-51.
  - **Shared License** (ASA only.)—Use the SSL VPN Shared License page to configure your SSL VPN Shared License. For more information, see Configuring SSL VPN Shared Licenses (ASA 8.2+), page 31-74.

  - **SSL VPN** (IOS devices only.)—The SSL VPN policy table lists all of the contexts that define the virtual configurations of the SSL VPN. Each context has a gateway, domain or virtual hostname, and user group policies. For more information, see Configuring an SSL VPN Policy (IOS), page 33-14.

Discovering Remote Access VPN Policies

Security Manager allows you to import the configurations of remote access IPSec VPN policies during policy discovery. You can also discover SSL VPN policies on ASA devices, but not on IOS devices. To discover remote access VPN policies, select the **RA VPN Policies** option in the Discover Device settings when adding the device to the inventory or when discovering policies on a device already in the inventory. For more information on adding devices or discovering policies, see the following topics:

- Adding Devices to the Device Inventory, page 3-7
- Discovering Policies on Devices Already in Security Manager, page 5-15

You can discover configurations on devices that are already deployed in your remote access VPN network, so that Security Manager can manage them. These configurations are imported into Security Manager as remote access VPN policies. Remote access VPN policy discovery can be performed by importing the configuration of a live device or by importing a configuration file. However, SSL VPN policies that refer to files in flash storage cannot be discovered from configuration files, therefore, we recommend that you do not discover SSL VPNs from configuration files.

When you initiate policy discovery on a device in a remote access VPN, the system analyzes the configuration on the device and then translates this configuration into Security Manager policies so that the device can be managed. Warnings are displayed if the imported configuration completes only a partial policy definition. If additional settings are required, you must go to the relevant page in the Security Manager interface to complete the policy definition. You can also rediscover the configurations of devices that are already managed with Security Manager.

When discovering SSL VPN policies, files residing in flash storage that are referenced in SSL VPN policies are copied to the Security Manager server to be stored in the /csm directory on the target device when policies are deployed from Security Manager. If the flash storage contains files that you want to use, but they are not referenced by an SSL VPN policy, either configure commands that refer to them or manually copy them to the Security Manager server. Policy discovery fails if an SSL VPN policy on the device refers to a file that has been deleted from flash; in this case, either fix the configuration directly before discovering the device, or deselect the **RA VPN Policies** option when adding the device and create the desired SSL VPN configuration in Security Manager.

Tips

- You should perform deployment immediately after you discover the policies on a device before you make any changes to policies or unassign policies from the device; otherwise, the changes that you configure in Security Manager might not be deployed to the device.
- For ASA and PIX 7.0+ devices, the default connection profiles and group policy are discovered and added to the Connection Profiles and Group Policies policy. You can modify these default profiles and group, but you cannot delete them:
  - DefaultRAGroup—The default connection profile for remote access IPsec VPNs.
  - DefaultWEBVPNGroup—The default connection profile for SSL VPNs. This connection profile is discovered only for ASA 8.0+ devices.

– DfltGrpPolicy—The default group policy, which is used by the default connection profiles. When discovered, Security Manager uses the name <*device_display_name*>DfltGrpPolicy. However, when you deploy configurations, the device display name is stripped off and DfltGrpPolicy is used.

  This naming convention is necessary because group policies are modeled as shared policy objects, and you might have modified the default group policy differently on your devices. However, the naming convention does not prevent you from using shared policies that incorporate the default group policy; the device display name is stripped from the object name regardless of the device to which it is assigned. For example, if you use the object 10.100.10.1DfltGrpPolicy with device 10.200.11.1, Security Manager still uses "DfltGrpPolicy" in the configuration.

Note Although these default connection profiles use the DfltCustomization object for SSL VPN portal customization, Security Manager does not discover it. To modify DfltCustomization, you must do so directly on the device. However, you can simply create your own customization object and specify it in the default connection profile to use non-default settings.

Related Topics

- Discovering Policies, page 5-12
- Site-To-Site VPN Discovery, page 25-20
- VPN Discovery Rules, page 25-22

Using the Remote Access VPN Configuration Wizard

You can use the Remote Access VPN Configuration wizard to create the policies required to configure a basic IPsec or SSL VPN. The wizard provides simplified options to configure the basic settings. Thus, after using the wizard, you might need to configure additional settings in the individual remote access VPN policies.

Tip The wizard never creates a valid IKEv2 IPSec VPN. You must always configure additional policies to complete an IKEv2 configuration.

Note For Remote Access VPN Multi-Context Mode, only Remote Access SSL VPN is supported for ASA devices running the software version 9.5(2) or later.

Depending on the device type and VPN type (IPSec or SSL), the wizard takes you through the steps to configure a basic remote access VPN.

To access the Remote Access Configuration wizard:

1. In Device view, select the device to configure as your remote access server from the Device selector.
2. Select **Remote Access VPN > Configuration Wizard** from the Policy selector.
3. Select the radio button corresponding to the type of remote access VPN you want to create: **Remote Access SSL VPN** or **Remote Access IPSec VPN**.
4. Click **Remote Access Configuration Wizard** to open the appropriate wizard.

For detailed information on how to use each version of the wizard, see the following topics:

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32
- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices), page 30-36

Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices)

This procedure describes how to create or edit SSL VPNs on ASA devices using the Remote Access SSL VPN Configuration Wizard.

Related Topics

- Understanding Remote Access SSL VPNs, page 30-2
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8

Step 1 In Device view, select the desired ASA device.

Step 2 From the Policy selector, select **Remote Access VPN > Configuration Wizard**.

Step 3 Select the **Remote Access SSL VPN** radio button.

Step 4 Click **Remote Access Configuration Wizard**. The Access page opens. For a description of the elements on this page, see SSL VPN Configuration Wizard—Access Page (ASA), page 30-16.

Step 5 Specify the interfaces on which you want to enable the SSL VPN connections. Click **Select** to select an interface or an interface role object that identifies the interfaces.

Step 6 Specify the port number you want to use for the SSL VPN sessions. Enter the port number or the name of a port list object that defines the number, or click **Select** to select the object or to create a new object.

The default port is 443, for HTTPS traffic. The port number can be 443, or within the range of 1024-65535. If you change the port number, all current SSL VPN connections terminate, and current users must reconnect.

Note If HTTP port redirection is enabled, the default HTTP port number is 80.

Step 7 To allow users to select a tunnel group from a list of tunnel group connection profiles configured on the device at login, select the **Allow Users to Select Connection Profile in Portal Page** option.

Step 8 To allow users to use the AnyConnect VPN client to connect to the SSL VPN, select the **Enable AnyConnect Access** check box.

Step 9 Click **Next**. The Connection Profile page opens. For a description of the elements on this page, see SSL VPN Configuration Wizard—Connection Profile Page (ASA), page 30-17.

Step 10 In **Connection Profile Name**, enter the name of the connection profile. This is the name of the tunnel group, and will appear in the Remote Access VPN > Connection Profiles policy. For more information about the connection profile policy, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.

Step 11 On the Connection Profile page, configure these options that will later appear in the General tab of the connection profile (see General Tab (Connection Profiles), page 31-10):

- **Group Policy**—Enter the name of the ASA Group Policy policy object that will be the default group for the connection profile, or click **Select** to select the object. If the required object does not yet exist, click **Select**, then click the **Create (+)** button in the ASA User Groups Selector dialog box, which opens a wizard to guide you through the creation process as described in Creating User Groups with the Create Group Policy Wizard, page 30-20.

  For more information about ASA Group Policies objects, see ASA Group Policies Dialog Box, page 34-1.

- **Group Policies**—This table lists all group policies currently used on the device, whether for SSL or IPsec VPNs. You can click **Edit** to add other group policies.

- **Global IP Address Pool**—Enter the address pools from which IP addresses are assigned. The server uses these address pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools.

  Specify the pools as address ranges or network/host objects that contain address ranges, in the format *Start_Address-End_Address*, for example, 10.100.10.2-10.100.10.254. Click **Select** to select network/host objects or to create new objects.

Step 12 On the Connection Profile page, configure these options that will later appear in the SSL VPN tab of the connection profile (see SSL Tab (Connection Profiles), page 31-22):

- **Portal Page Customization**—The name of the SSL VPN Customization policy object that defines the default portal page for the VPN. Click **Select** to select the object or to create a new object.

  Note You can set up different login windows for different groups by using a combination of customization profiles and tunnel groups. For example, assuming that you had created a customization profile called salesgui, you can create an SSL VPN tunnel group called sales that uses that customization profile.

- **Connection URL**—The URL of the connection profile. This URL provides users with direct access to the customized portal page. Select a protocol (**http** or **https**) from the list, and specify the URL including the name of the connection profile, in the field provided.

  The URL is made up of the host name or IP address of the ASA device and port number, and the alias used to identify the SSL VPN connection profile.

  Note If you do not specify a URL, you can access the portal page by entering the portal page URL, and then selecting the connection profile alias from a list of configured connection profile aliases configured on the device. See SSL VPN Configuration Wizard—Access Page (ASA), page 30-16.

Step 13 On the Connection Profile page, configure the AAA options for authentication, authorization, and accounting, and secondary authentication, which will later appear on the AAA and Secondary AAA tab of the connection profile (see AAA Tab (Connection Profiles), page 31-13 and Secondary AAA Tab (Connection Profiles), page 31-17).

Step 14 Click **Finish** to save your changes.

SSL VPN Configuration Wizard—Access Page (ASA)

Use the Access page of the SSL VPN Configuration Wizard to configure the security appliance interfaces for SSL VPN sessions. After you complete the wizard, you can later edit these settings in the SSL VPN Access policy; see SSL VPN Access Policy Page, page 31-45.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access SSL VPN on an ASA device (see Using the Remote Access VPN Configuration Wizard, page 30-13). The Access page is the first page that appears.

Related Topics

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 30-2 SSL VPN Wizard—Access Page (ASA)

Element	Description
Interfaces to Enable SSL VPN Service	The interfaces or interface roles that identify the interfaces on which you want to enable SSL VPN connections. Click **Select** to select interfaces or interface roles, or to create new interface roles.
Port Number	The port number to use for the SSL VPN sessions. Enter a port number or port list object name, or click **Select** to select an object that defines the port, or to create a new object. The default port is 443, for HTTPS traffic. The port number can be 443, or within the range of 1024-65535. If you change the port number, all current SSL VPN connections terminate, and current users must reconnect. **Note** If HTTP port redirection is enabled, the default HTTP port number is 80.
Portal Page URLs	The URLs that users would use to connect to the VPN. The URLs are displayed after you specify the interfaces and port number.
Allow Users to Select Connection Profile in Portal Page	Whether to present a list of configured connection profiles (tunnel groups) from which the user can select the appropriate profile when the user logs in (for example, in the SSL VPN portal page). If you do not select this option, the user cannot select a profile and must use the default profile for the connection.

Table 30-2 SSL VPN Wizard—Access Page (ASA) (continued)

Element	Description
Enable AnyConnect Access	Whether to allow the user to use the AnyConnect VPN client to make an SSL or IKEv2 IPSec VPN connection. The option is selected by default. For details about AnyConnect VPN clients, see Understanding SSL VPN AnyConnect Client Settings, page 31-62. **Note** To enable AnyConnect Essentials, go to Remote Access VPN > SSL VPN > Access. For details, see Configuring an Access Policy, page 31-49.

SSL VPN Configuration Wizard—Connection Profile Page (ASA)

Use the Connection Profile page in the SSL VPN Configuration wizard to configure the tunnel group policies on your security appliance. You can specify a name for the tunnel connection profile policy that you are adding, select the user group policy, specify address pools for this policy, and specify authentication server group settings.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access SSL VPN on an ASA device (see Using the Remote Access VPN Configuration Wizard, page 30-13); then click **Next** until you reach this page.

Related Topics

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- ASA Group Policies Dialog Box, page 34-1
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78
- Understanding Networks/Hosts Objects, page 6-80
- Understanding AAA Server and Server Group Objects, page 6-27

Field Reference

Table 30-3 SSL VPN Configuration Wizard, Connection Profile Page (ASA)

Element	Description
Connection Profile Name	The name of the connection profile (tunnel group).

Table 30-3 SSL VPN Configuration Wizard, Connection Profile Page (ASA) (continued)

Element	Description
Group Policy	Default ASA user group associated with the device. Enter an ASA user group policy or click **Select** to select one from a list or to create a new one. If required, the name of the ASA group policy object that defines the default user group associated with the connection profile. A group policy is a collection of user-oriented attribute/value pairs stored either internally on the device or externally on a RADIUS/LDAP server. Click **Select** to select an existing object or to create a new one. If you click the **Create (+)** button in the group policy selection dialog box, you are guided through the group creation process using a wizard, as explained in Creating User Groups with the Create Group Policy Wizard, page 30-20.
Full Tunnel	A read-only field that indicates whether full tunnel access mode is configured for the object selected in the Group Policy field.
Group Policies	The names of all ASA user group policies that are configured for the device, even those that are configured for IPSec VPN connections only. The contents of this table is identical to the contents of the Remote Access VPN > Group Policies policy. The table shows whether full tunnel access mode is enabled or disabled for each group policy. You can change the list by clicking **Edit**. This opens a dialog box where you can select additional group policies, or deselect currently selected policies (do not deselect policies that are used by other connection profiles). You can also create new group policies (click the **Create (+)** button below the available group policies list) or edit the group policy object by selecting it and clicking the **Edit (pencil)** button below either list. If you create a new group policy, the Create Group Policy wizard is used to guide you through the process. See Creating User Groups with the Create Group Policy Wizard, page 30-20.
Portal Page Customization	The name of the SSL VPN Customization policy object that defines the default portal page for the VPN. This profile defines the appearance of the portal page that allows the remote user access to all resources available on the SSL VPN. Click **Select** to select the object or to create a new object.
Connection URL	The URL of the connection profile. This URL provides users with direct access to the customized portal page. Select a protocol (**http** or **https**) from the list and specify the URL, including host name or IP address of the ASA device and port number and the alias used to identify the SSL VPN connection profile. **Note** If you do not specify a URL, you can access the portal page by entering the portal page URL, and then selecting the connection profile alias from a list of configured connection profile aliases configured on the device. See SSL VPN Configuration Wizard—Access Page (ASA), page 30-16.

Table 30-3 SSL VPN Configuration Wizard, Connection Profile Page (ASA) (continued)

Element	Description
Global IPv4 Address Pool	The address pools from which IPv4 addresses will be assigned to clients if no pool is specified for the interface to which the client connects. Address pools are entered as a range of addresses, such as 10.100.12.2-10.100.12.254. The server uses these pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools. Enter the address pool ranges or the names of network/host objects that define these pools. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.
Global IPv6 Address Pool	The address pools from which IPv6 addresses will be assigned to clients if no pool is specified for the interface to which the client connects. Address pools are entered as a range of addresses, such as 2001:db8::1-2001:db8::2:1. The server uses these pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools. . Enter the address pool ranges or the names of network/host objects that define these pools. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.
Authentication Server Group	The name of the authentication server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Use LOCAL if Server Group Fails	Whether to fall back to the local database for authentication if the selected authentication server group fails.
Authorization Server Group	The name of the authorization server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Accounting Server Group	The name of the accounting server group. Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Secondary Authentication	Whether to enable double authentication, which prompts the user for two sets of credentials (username and password) before completing the remote access VPN connection. • **Enable Secondary Authentication**—Select this option to require double authentication. • **Authentication Server Group**—The name of the authentication server group (LOCAL if the tunnel group is configured on the local device) to be used with the second set of credentials. Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object. • **Use LOCAL if Server Group Fails**—Whether to fall back to the local database for authentication if the selected authentication server group fails.

Creating User Groups with the Create Group Policy Wizard

When you are using the Remote Access SSL VPN Configuration wizard to create an SSL VPN on ASA or IOS devices, you can create new ASA group policy or IOS user group objects using a wizard. The wizard lets you configure select elements of the group, so you might need to edit the object after creating it to configure additional settings.

The Create Group Policy wizard is available only through the Remote Access SSL VPN Configuration wizard. For an explanation of how to start and use the wizard, see the following topics:

The following procedure assumes that you are already in the Remote Access SSL VPN Configuration wizard, as described in the following topics:

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32

Related Topics

- SSL VPN Access Modes, page 30-4

Step 1 When using the Remote Access VPN Configuration wizard for SSL VPNs, proceed to the page where you select group policies. On this page, you can open the selection page for user groups by doing the following:

- ASA devices—On the Connection Profile page of the wizard, click **Select** next to the Group Policy field, or click **Edit** next to the Group Policies table.
- IOS devices—On the Gateway and Context page of the wizard, click **Edit** next to the Group Policies table.

Step 2 In the Group Policy Selector dialog box, click the **Create (+)** button below the list of available group policies to start the Create Group Policy wizard. The wizard starts at the Group Policy page.

You can also do the following on the Group Policy Selector:

- Select existing groups and click >> to use them in the SSL VPN. When selecting a group for the default group on an ASA (the Group Policy field), you select the object simply by clicking it in the list.
- Select an existing group and click **Edit (pencil)** to change its properties.

Step 3 On the Group Policy page, configure the following options:

- **Name**—The name of the user group. Enter up to 128 characters, including uppercase and lowercase characters and most alphanumeric or symbol characters.
- **Access Method**—Select the required remote access method options, as follows:
  - **Full Tunnel**—To access to the corporate network completely over an SSL VPN tunnel. This is the recommended option.
  - **Clientless**—To access the internal or corporate network using a web browser on the client machine.
  - **Thin Client**—To download a Java applet that acts as a TCP proxy on the client machine.

Step 4 Click **Next**. The page that opens next depends on which access methods you selected. This procedure assumes that you selected all methods, in which case the Full Client page opens.

Step 5 On the Full Client page, select whether to restrict access to full tunnel only or to allow other methods of access if the full client download fails. Also, specify DNS and WINS server information, and configure split tunneling if you want to allow it. For an explanation of the options, see Create Group Policy Wizard—Full Tunnel Page, page 30-21.

Step 6 Click **Next**. The Clientless and Thin Client page opens.

Step 7 On the Clientless and Thin Client page, configure these access modes. For an explanation of the options, see Create Group Policy Wizard—Clientless and Thin Client Access Modes Page, page 30-24.

Step 8 Click **Finish** to create the group policy object.

Step 9 When you complete the wizard, the group policy is added to the available groups list, but it is not selected (unless you are configuring the default group for an ASA). To select it, highlight it in the available groups list and click **>>** to move it to the selected groups list.

Note To specify a user group as the default user group, select it and click **Set As Default**. This option is only available for IOS routers.

Step 10 Click **OK** in the Group Policy Selector page to save your changes and return to the Remote Access SSL VPN Configuration wizard.

Create Group Policy Wizard—Full Tunnel Page

Note This page is available only if you selected the **Full Client** option in the Group Policy of the Create Group Policy wizard.

In this page, you can configure the mode used to access the corporate network.

Navigation Path

For information on starting the Create Group Policy wizard, see Creating User Groups with the Create Group Policy Wizard, page 30-20.

Field Reference

Table 30-4 Create User Group Wizard—Full Tunnel Page

Element	Description
Mode	The access modes to allow in the SSL VPN. Select one of the following: • **Use Other Access Modes if SSL VPN Client Download Fails**—To allow the remote client to use clientless or thin client access modes if the download of the VPN client fails. • **Full Tunnel Only**—Prohibit clientless or thin client access. The user must have the full client installed and functional to connect to the VPN. Ensure that you configure the full client images on the device. For ASA devices, use the Client Settings tab of the SSL VPN > Other Settings policy; see Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64. For IOS devices, the client is managed using a FlexConfig policy; see Predefined FlexConfig Policy Objects, page 7-19.
Client IP Address Pools (IOS device only.)	The IP address ranges of the address pool that full tunnel clients will draw from when they log on. The address pool must be in the same subnet as one of the device's interface IP addresses. Enter the address range separating the first and last IP address with a hyphen, for example, 10.100.10.2-10.100.10.255. If you enter a single address, the pool has just one address. Do not enter subnet designations. You can also enter the name of a network/host policy object that defines the range, or click **Select** to select the object from a list or to create a new object. Separate multiple ranges with commas.
Primary IPv4 DNS Server	The IPv4 address of the primary DNS server for the group. Enter the IPv4 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary IPv4 DNS Server	The IPv4 address of the secondary DNS server for the group. Enter the IPv4 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Primary IPv6 DNS Server	The IPv6 address of the primary DNS server for the group. Enter the IPv6 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary IPv6 DNS Server	The IPv6 address of the secondary DNS server for the group. Enter the IPv6 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Default DNS Domain	The domain name of the DNS server to be used for Full Client SSL VPN connections.
Primary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.

Table 30-4 Create User Group Wizard—Full Tunnel Page (continued)

Element	Description
Split Tunnel Option	Whether to allow split tunneling for IPv4 traffic and if so, which traffic should be secured or transmitted unencrypted across the public network: • Disabled—(Default) No IPv4 traffic goes in the clear or to any other destination than the gateway. Remote users reach networks through the corporate network and do not have access to local networks. • Tunnel Specified Traffic—Tunnel all IPv4 traffic from or to the addresses listed in the **Networks** or **Destinations** field. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. • Exclude Specified Traffic—IPv4 traffic goes in the clear from and to the addresses listed in the **Networks** or **Destinations** field. This is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel.
IPv6 Split Tunnel Option	Whether to allow split tunneling for IPv6 traffic and if so, which traffic should be secured or transmitted unencrypted across the public network: • Disabled—(Default) No IPv6 traffic goes in the clear or to any other destination than the gateway. Remote users reach networks through the corporate network and do not have access to local networks. • Tunnel Specified Traffic—Tunnel all traffic from or to the addresses listed in the **Networks** or **Destinations** field. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. • Exclude Specified Traffic—Traffic goes in the clear from and to the addresses listed in the **Networks** or **Destinations** field. This is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel.
Networks (ASA device only.)	If you select Tunnel Specified Traffic or Exclude Specified traffic in the Split Tunnel Option, enter the name of the ACL object that defines the traffic to be tunneled or excluded. Click **Select** to select the object or to create a new object.
Destinations (IOS device only.)	If you select Tunnel Specified Traffic or Exclude Specified traffic in the Split Tunnel Option, specify the IP addresses that define the traffic to be tunneled or excluded. Enter network addresses such as 10.100.10.0/24 or host addresses such as 10.100.10.12. You can also enter the name of a network/host policy object, or click **Select** to select the object from a list or to create a new object. Separate multiple addresses with commas.

Table 30-4 Create User Group Wizard—Full Tunnel Page (continued)

Element	Description
Exclude Local LANs (IOS device only.)	Whether to exclude local LANs from the encrypted tunnel. This option is available only if you selected the **Exclude Specified Traffic** split tunnel option. By selecting this option, you do not have to enter local LAN addresses into the destinations field to allow users to communicate with systems (such as printers) that are attached to their LAN. When selected, this attribute disallows a non split-tunneling connection to access the local subnetwork at the same time as the client.
Split DNS Names	A list of domain names to be resolved through the split tunnel to the private network. All other names are resolved using the public DNS server. Enter up to 10 entries in the list of domains, separated by commas. The entire string can be no longer than 255 characters.

Create Group Policy Wizard—Clientless and Thin Client Access Modes Page

In the Clientless and Thin Client page of the Create Group Policy wizard, you can configure the Clientless and Thin Client modes to be used for accessing the corporate network in your SSL VPN.

Note This page is only available if you selected the **Clientless** or **Thin Client** options in step 1 of the Create Group Policy wizard.

Navigation Path

For information on starting the Create Group Policy wizard, see Creating User Groups with the Create Group Policy Wizard, page 30-20.

Related Topics

- SSL VPN Access Modes, page 30-4
- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Add or Edit Port Forwarding List Dialog Boxes, page 34-40

Field Reference

Table 30-5 Create User Group Wizard—Clientless and Thin Client Page

Element	Description
Clientless—Appears only if you selected **Clientless** in step 1 of the wizard.	
Portal Page Websites	The name of the SSL VPN bookmarks policy object that includes the web site URLs to display on the portal page. These web sites help users access desired resources. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Allow Users to Enter Websites	Whether to allow the remote user to enter web site URLs directly into the browser. If you do not select this option, the user can access only those URLs included on the portal.

Table 30-5 Create User Group Wizard—Clientless and Thin Client Page (continued)

Element	Description
Thin Client—Appears only if you selected **Thin Client** in step 1 of the wizard.	
Port Forwarding List	The name of the port forwarding list policy object assigned to this group. Port forwarding lists contain the set of applications that users of clientless SSL VPN sessions can access over forwarded TCP ports. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Port Forwarding Applet Name (ASA device only.)	The application name or short description to display on the Port Forwarding Java applet screen on the portal, up to 64 characters. This is the name of the applet users will download to act as a TCP proxy on the client machine for the services configured on the SSL VPN gateway.
Download Port Forwarding Applet on Client Login	Whether the port forwarding Java applet should be automatically downloaded to the client when a user logs into the SSL VPN. If you do not automatically download the applet, users must download it manually after login.

Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

This procedure describes how to create IPSec VPNs on ASA or PIX 7.0+ devices using the Remote Access VPN Configuration Wizard.

Tip The wizard allows you to select shared policies to use in the VPN on the Defaults page (the final step of the wizard). If you want to use this feature, you must first ensure that all required shared policies are configured and submitted to the database. For information on configuring shared policies and VPN policy defaults, see Understanding and Configuring VPN Default Policies, page 25-12.

Related Topics

- Understanding Remote Access IPSec VPNs, page 30-2
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8

Step 1 In Device view, select the desired ASA or PIX 7.0+ device.

Step 2 From the Policy selector, select **Remote Access VPN > Configuration Wizard**.

Step 3 Select the **Remote Access IPSec VPN** radio button.

Step 4 Click **Remote Access Configuration Wizard**. The Connection Profile page opens. For a description of the options that appear on this page, see Remote Access VPN Configuration Wizard—IPSec VPN Connection Profile Page (ASA), page 30-28.

Step 5 On the Connection Profile page, configure these basic options:

- **Connection Profile name**—Enter the name of the connection profile. This is the name of the tunnel group, and will appear in the Remote Access VPN > Connection Profiles policy. For more information about the connection profile policy, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.
- **IKE Versions**—Select the IKE versions to use during IKE negotiations between the VPN server and the remote users, version 1, 2 or both. IKEv2 is supported on ASA Software release 8.4(1)+ only.

Step 6 On the Connection Profile page, configure these options that will later appear in the General tab of the connection profile (see General Tab (Connection Profiles), page 31-10):

- **Group Policy**—Enter the name of the ASA Group Policy policy object that will be the default group for the connection profile, or click **Select** to select the object. If the required object does not yet exist, click **Select**, then click the **Create (+)** button in the ASA User Groups Selector dialog box to open the dialog boxes that are used to create these objects.

  When creating a new group policy object, you must select the same IKE versions that you select in the Connection Profile page of the wizard. These options are on the Technology page of the Add ASA Group Policies dialog box: Easy VPN/IPSec IKEv1 and Easy VPN/IPSec IKEv2.

  For more information about ASA Group Policies objects, see ASA Group Policies Dialog Box, page 34-1.
- **Global IP Address Pool**—Enter the address pools from which IP addresses are assigned. The server uses these address pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools.

  Specify the pools as address ranges or network/host objects that contain address ranges, in the format *Start_Address-End_Address*, for example, 10.100.10.2-10.100.10.254. Click **Select** to select network/host objects or to create new objects.

Step 7 On the Connection Profile page, configure the AAA options for authentication, authorization, and accounting, which will later appear on the AAA tab of the connection profile (see AAA Tab (Connection Profiles), page 31-13).

Step 8 Click **Next** to move to the IPsec Settings page.

Step 9 On the IPSec Settings page, configure the options for IPSec, which will later appear on the IPSec tab of the connection profile (see IPSec Tab (Connection Profiles), page 31-19). Note that some of these settings apply to IKEv1 only.

- **Preshared Key**, **Confirm**—Enter the IKEv1 preshared key for the tunnel group in each field. The maximum length of a preshared key is 127 characters.

  You cannot configure a preshared key for remote access IKEv2 IPsec VPNs.
- **Trustpoint Name**—Enter the name of the PKI enrollment policy object that defines the trustpoint name, if any trustpoints are configured, for an IKEv1 connection. A trustpoint represents a Certificate Authority (CA)/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate. Click **Select** to select the object from a list or to create a new object.

  For IKEv2, the trustpoint name is not configured here, but on the IKEv2 Settings tab of the Global Settings policy. The configuration is explained later in this procedure.
- The other options (other than the client table) apply to both IKEv1 and IKEv2. Change the settings if you need non-default behavior. For an explanation of the options, including the client software update table, see Remote Access VPN Configuration Wizard—IPSec Settings Page (ASA), page 30-30.

Step 10 Click **Next** to move to the VPN Defaults page.

Step 11 On the Defaults page, select the additional shared policies that you want to assign to the VPN. Initially, the policies listed are those chosen on the Security Manager Administration VPN Defaults page.

For more information about selecting these policies, see Remote Access VPN Configuration Wizard—Defaults Page, page 30-31.

Step 12 Click **Finish** to save your changes.

Because the wizard does not configure all possible options, inspect the policies created and configure any additional options that you want to implement.

The remaining steps are required if you selected IKE version 2 as a supported IKE version, or if you specified an IPsec trustpoint.

Step 13 (IKEv2 Optional.) Configure group aliases and double authentication if required:

a. Select the Connection Profiles policy.

b. Select the connection profile you configured in the wizard, and click the **Edit Row (pencil)** button to open the Connection Profiles dialog box.

- If you want to configure double authentication, select the Secondary AAA tab and configure the required settings. For more information, see Secondary AAA Tab (Connection Profiles), page 31-17.
- If you want to configure aliases for the profile, which helps users select the correct profile during login, select the SSL tab and configure the alias table. For more information, see SSL Tab (Connection Profiles), page 31-22.
- There are several additional connection profile settings that are not configured in the wizard. Examine the tabs in the Connection Profile dialog box to determine if additional changes are required.

c. Click **OK** in the Connection Profiles dialog box to save your changes.

Step 14 (IKEv2 Requirement.) Select the **Remote Access VPN > SSL VPN > Access** policy and configure at least the following. For detailed information about configuring an Access policy, see Understanding SSL VPN Access Policies (ASA), page 31-44.

- Add the remote access VPN interface to the access interfaces table.
- Select **Allow Users to Select Connection Profile in Portal Page**.
- Select **Enable AnyConnect Access**.

Step 15 (IKEv2 Requirement.) Select the **Remote Access VPN > SSL VPN > Other Settings** policy, and click the **Client Settings** tab.

In the AnyConnect Client Image table, add an AnyConnect 3.0 or higher client image, one that supports IKEv2 negotiations.

For more information on configuring client images, see Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64.

Step 16 (IKEv2 Requirement.) Select the **Remote Access VPN > Global Settings** policy, and click the IKEv2 Settings tab.

At minimum, configure the **RA Trustpoint** for remote access IKEv2 authentication. Enter the name of the PKI enrollment object that identifies the certificate authority (CA) server or click **Select** to select the object or to create a new one.

For more information on configuring IKEv2 global settings, see Configuring VPN Global IKEv2 Settings, page 26-37.

Step 17 (IKEv1, IKEv2 Requirement.) Select the **Remote Access VPN > Public Key Infrastructure** policy and ensure that the following PKI enrollment objects are selected:

- (IKEv1) The object specified on the IPSec tab of the connection profile, if a trustpoint is configured.
- (IKEv2) The object specified on the IKEv2 Settings tab of the Global Settings policy.

Note In the wizard, you might have applied a shared Public Key Infrastructure policy that already specifies these objects.

Step 18 (IKEv2 Optional.) IKEv2 connections require the use of the AnyConnect 3.0+ client. The AnyConnect client might need to download files, such as software upgrades, profiles, localization and customization files, CSD, SCEP, and so forth. The wizard does not enable these types of download.

To enable AnyConnect file downloads:

a. Select **Remote Access VPN > IPSec VPN > IPSec Proposal**.

b. Select the IPSec proposal created by the wizard, and click **Edit Row (pencil)** to open the IPSec Proposal Editor. For information about the various options, see IPsec Proposal Editor (ASA, PIX 7.0+ Devices), page 31-41.

c. Select the **Enable Client Services** option, and enter a port number if you do not want to use the default port 443. (You can use the same port number used for SSL VPN or other SSL uses.)

d. Click **OK** to save your changes.

Remote Access VPN Configuration Wizard—IPSec VPN Connection Profile Page (ASA)

Use the Connection Profile page of the Remote Access VPN Configuration wizard to configure the connection profile policies on your security appliance for a remote access IPSec VPN. You can specify a name for the connection profile policy that you are adding, select the IKE versions to allow during IKE negotiations, select the user group policy, specify address pools for this policy, and specify authentication, authorization, and accounting server group settings.

For more information about using the wizard to configure remote access IPsec VPNs on ASA, see Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access IPsec VPN on an ASA or PIX 7.0+ device (see Using the Remote Access VPN Configuration Wizard, page 30-13). The IPSec Connection Profile page is the first page that appears.

Field Reference

Table 30-6 Remote Access VPN Configuration Wizard, IPSec Connection Profile Page (ASA)

Element	Description
Connection Profile Name	The name of the connection profile (tunnel group).

Table 30-6 Remote Access VPN Configuration Wizard, IPSec Connection Profile Page (ASA)

Element	Description
IKE Versions	The IKE versions to use during IKE negotiations between the VPN server and the remote users. IKEv2 is supported on ASA Software release 8.4(1)+ only; you cannot change the option selection on other types of device. Select **IKE Version 1**, **IKE Version 2**, or **Both** (to allow either version). IKEv2 connections are allowed using Anyconnect clients only.
Group Policy	If required, the name of the ASA group policy object that defines the default user group associated with the connection profile. A group policy is a collection of user-oriented attribute/value pairs stored either internally on the device or externally on a RADIUS/LDAP server. Click **Select** to select an existing object or to create a new one. **Tip** If you enable IKEv2 for this VPN, there are special considerations for the group policy you choose. For detailed information, see Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25.
Global IP Address Pool	The address pools from which IP addresses will be assigned to clients if no pool is specified for the interface to which the client connects. Address pools are entered as a range of addresses, such as 10.100.12.2-10.100.12.254. The server uses these pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools. Enter the address pool ranges or the names of network/host objects that define these pools. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.
Authentication Server Group	The name of the authentication server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Use LOCAL if Server Group Fails	Whether to fall back to the local database for authentication if the selected authentication server group fails.
Authorization Server Group	The name of the authorization server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Accounting Server Group	The name of the accounting server group. Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.

Remote Access VPN Configuration Wizard—IPSec Settings Page (ASA)

Use the IPSec Settings page of the Remote Access VPN Configuration wizard to configure the IPSec settings on your security appliance for a remote access IPSec VPN. Some of these settings apply to IKE version 1 (IKEv1) only; if you are configuring an IKEv2-only VPN, these fields are greyed and unconfigurable.

For more information about using the wizard to configure remote access IPsec VPNs on ASA, see Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access IPsec VPN on an ASA or PIX 7.0+ device (see Using the Remote Access VPN Configuration Wizard, page 30-13); then click **Next** until you reach this page.

Field Reference

Table 30-7 Remote Access VPN Configuration Wizard, IPSec VPN Wizard—IPSec Settings (ASA)

Element	Description
Preshared Key (IKEv1 only.)	The preshared key for the connection profile. The maximum length of a preshared key is 127 characters. Enter the key again in the Confirm field. **Tip** You cannot configure preshared keys for IKEv2 remote access VPNs.
Trustpoint Name (IKEv1 only.)	The name of the PKI enrollment policy object that defines the trustpoint name if any trustpoints are configured for IKEv1 connections. A trustpoint represents a Certificate Authority (CA)/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate. Click **Select** to select the object from a list or to create a new object. **Tip** This trustpoint is used for IKEv1 negotiations only. To configure the global trustpoint for IKEv2 negotiations, use the IKEv2 Settings tab of the Global Settings policy; see Configuring VPN Global IKEv2 Settings, page 26-37.
IKE Peer ID Validation	Select whether IKE peer ID validation is ignored (Do not check), required, or checked only if supported by a certificate. During IKE negotiations, peers must identify themselves to one another.
Enable Sending Certificate Chain	Whether to enable the sending of the certificate chain for authorization. A certificate chain includes the root CA certificate, identity certificate, and key pair.
Enable Password Update with RADIUS Authentication	When selected, enables passwords to be updated with the RADIUS authentication protocol. Whether to enable passwords to be updated with the RADIUS authentication protocol. For more information, see Supported AAA Server Types, page 6-28.

Table 30-7 Remote Access VPN Configuration Wizard, IPSec VPN Wizard—IPSec Settings (ASA)

Element	Description
ISAKMP Keepalive	Whether to monitor ISAKMP keepalive. If you select the **Monitor Keepalive** option, you can configure IKE keepalive as the default failover and routing mechanism. Enter the following parameters: • **Confidence Interval**—The number of seconds that a device waits between sending IKE keepalive packets. • **Retry Interval**—The number of seconds a device waits between attempts to establish an IKE connection with the remote peer. The default is 2 seconds. For more information, see Configuring VPN Global ISAKMP/IPsec Settings, page 26-33.
Client Software Update table (IKEv1 only.)	The VPN client revision level and URLs for client platforms. You can configure different revision levels for All Windows Platforms, Windows 95/98/ME, Windows NT4.0/2000/XP, or the VPN3002 Hardware Client. To configure the client for a platform, select it, click the **Edit Row** button, and fill in the IPSec Client Software Update Dialog Box, page 31-22.

Remote Access VPN Configuration Wizard—Defaults Page

Use the Defaults page of the Remote Access VPN Configuration wizard to select the shared policies to assign to the remote access IPSec VPN. Initially, the policies selected are those configured in the Security Manager Administration VPN Defaults for remote access VPNs. For information on how to configure these defaults, see Understanding and Configuring VPN Default Policies, page 25-12.

Required policies must always have a policy selected. If "Factory Default" is shown, then the policy applied is not a shared policy but default policy settings supplied by Security Manager. If you can select the empty option, the policy is optional and you need to configure it only if you want the associated features.

When evaluating which policies to assign (if any), keep the following in mind:

- The drop-down lists for each policy type list the existing shared policies that you can select. You can select only shared policies that have been committed to the Security Manager database (and approved, if you are using Workflow mode with an approver). You cannot create a shared policy and use it before you submit it.
- To view the content of a policy, select it and click the **View Content** button. You are presented with a read-only view of the policy. Use this to help verify that you are selecting the desired policy.

Note If you try to select a default policy that is currently locked by another user, a message is displayed warning you of a lock problem. To bypass the lock, select a different policy or cancel the VPN creation until the lock is removed. For more information, see Understanding Policy Locking, page 5-8.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access IPsec VPN (see Using the Remote Access VPN Configuration Wizard, page 30-13) and click **Next** until you reach this page.

Related Topics

- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25
- Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices), page 30-36
- Overview of Remote Access VPN Policies, page 30-9

Field Reference

Table 30-8 Remote Access VPN Configuration Wizard, Defaults Page

Element	Description
ASA Cluster Load Balance	Defines load balancing for an ASA device in your remote access VPN.
High Availability	Defines a High Availability (HA) policy on a Cisco IOS router in a remote access VPN.
Certificate to Connection Profile Map Policy	(IKEv1 only.) Defines the certificate to connection profile map options for an ASA device in your remote access VPN.
IKE Proposal	Defines the set of algorithms that two peers use to secure the IKE negotiation between them.
IPSec Proposal	Defines the crypto maps required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA.
Public Key Infrastructure	Defines the Public Key Infrastructure (PKI) policy used to generate PKI enrollment requests for PKI certificates and RSA keys.
VPN Global Settings	Defines global settings for IKE, IPsec, IKEv2, NAT, and fragmentation that apply to devices in your remote access VPN.

Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices)

This procedure describes how to create or edit SSL VPNs on IOS devices using the Remote Access SSL VPN Configuration Wizard.

Related Topics

- Understanding Remote Access SSL VPNs, page 30-2
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8

Step 1 In Device view, select the desired IOS device.

Step 2 From the Policy selector, select **Remote Access VPN > Configuration Wizard**.

Step 3 Select the **Remote Access SSL VPN** radio button.

Step 4 Click **Remote Access Configuration Wizard**. The Gateway and Context page opens. For a description of the elements on this page, see SSL VPN Configuration Wizard—Gateway and Context Page (IOS), page 30-33.

Step 5 Select the gateway to be used as a proxy for connections to the protected resources in your SSL VPN. Options are:

- **Use Existing Gateway**—Lets you use an existing gateway object. If you select this option, specify the name of the SSL VPN Gateway policy object that defines the gateway. Click **Select** to select the object or to create a new object.
- **Create Using IP Address**—Lets you configure a new gateway object using a reachable (public, static) IP address on the router. Enter the IP address.
- **Create Using Interface**—Lets you configure a new gateway using the public, static IP address of a router interface. Select the interface or interface role object.

If you elected to create a new gateway using an IP address or an interface:

- Specify a gateway name.
- Specify the number of the port that will carry the HTTPS traffic. The default is 443, unless HTTP port redirection is enabled, in which case the default HTTP port number is 80. If you want to use a different port, it must be between 1024 and 65535.

Step 6 Enter the name of the name of the context that defines the virtual configuration of the SSL VPN.

Step 7 Select the user groups that will be used in your SSL VPN policy. User groups define the resources available to users when connecting to an SSL VPN gateway. The table shows whether full client access is enabled for the group. Click **Edit** to select the desired groups, or to create new groups.

Step 8 Configure the AAA options for authentication, authentication domain, and accounting. For detailed information, see SSL VPN Configuration Wizard—Gateway and Context Page (IOS), page 30-33.

Step 9 Click **Next**. The Portal Page Customization page opens. For a description of the elements on this page, see SSL VPN Configuration Wizard—Portal Page Customization Page (IOS), page 30-35.

Step 10 On the Portal Customization page, configure the following options. The bottom of the page is a preview of what the portal page will look like based on your selections; use the preview to fine-tune your selections.

- **Title**—The name of the portal page, which appears at the top of the page.
- **Logo**—The graphic to show in the title area of the page: None, Default (the Cisco logo graphic), or Custom. If you select custom, click **Select** to select the graphic on the Security Manager server. You must copy the custom graphic to the server before you can use it in the portal customization.

  The source image file for the logo can be a GIF, JPG, or PNG file, with a file name of up to 255 characters, and up to 100 kilobytes in size.
- **Login Message**—The text that should appear above the login prompt.
- **Title and Text Colors**—The colors to use for the title and login area and the fonts.

Step 11 Click **Finish** to save your changes.

SSL VPN Configuration Wizard—Gateway and Context Page (IOS)

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

A gateway and context must be configured on a device before a remote user can access resources on a private network behind the SSL VPN. Use this step of the SSL VPN Configuration wizard to specify a gateway and context configuration, including information that will allow users to access a portal page.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access SSL VPN on an IOS device (see Using the Remote Access VPN Configuration Wizard, page 30-13). The Gateway and Context page is the first page that appears.

Related Topics

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32
- Add or Edit SSL VPN Gateway Dialog Box, page 34-64
- Understanding AAA Server and Server Group Objects, page 6-27

Field Reference

Table 30-9 SSL VPN Configuration Wizard, Gateway and Context Page

Element	Description
Gateway	The gateway to be used as a proxy for connections to the protected resources in your SSL VPN. Options are: • **Use Existing Gateway**—When selected, enables you to use an existing gateway for your SSL VPN. • **Create Using IP Address**—When selected, enables you to configure a new gateway using a reachable (public static) IP address on the router. • **Create Using Interface**—When selected, enables you to configure a new gateway using the public static IP address of the router interface.
Gateway Name	The name of the SSL VPN gateway policy object that defines the gateway: • If you selected Use Existing Gateway, click **Select** to select the object from a list or to create a new object. **Note** After selecting the gateway, the port number and digital certificate required to establish a secure connection are displayed in the relevant fields. • If you selected Create Using IP Address or Interface, enter the name of the object that you want to create (up to 128 characters).
IP Address	Available if you selected to create a gateway using an IP address. The IP address on the router that should be used as the gateway address.
Interface	Available if you selected to create a gateway using an interface. The name of the interface, or the interface role object that defines the interface, that should be used as the SSL VPN gateway. Click **Select** to select the interface or interface role, or to create a new interface role.

Table 30-9 SSL VPN Configuration Wizard, Gateway and Context Page (continued)

Element	Description
Port	The port number used for SSL VPN connections. The default is 443, unless HTTP port redirection is enabled, in which case the default HTTP port number is 80. If you enter a different number, it must be between 1024 and 65535. • If you selected Use Existing Gateway, this is a read-only field that shows the port number configured in the selected object. • If you selected Create Using IP Address or Interface, enter the port number or the name of a port list object that specifies the number, or click **Select** to select the port list object.
Trustpoint	The digital certificate required to establish the secure connection. A self-signed certificate is generated when an SSL VPN gateway is activated.
Context Name	The name of the context that defines the virtual configuration of the SSL VPN. **Note** To simplify the management of multiple context configurations, make the context name the same as the domain or virtual hostname.
Portal Page URL	The URL for the SSL VPN, which is filled in when you select (or define) a gateway object. Users connect to this URL to enter the VPN.
Group Policies	The user groups that will be used in your SSL VPN policy. User groups define the resources available to users when connecting to an SSL VPN gateway. The table shows whether full client access is enabled for the group. Click **Edit** to select the desired groups, or to create new groups.
Authentication Server Group	The authentication server groups. The list is in prioritized order. Authentication is attempted using the first group and proceeds through the list until the user is successfully authenticated or denied. Use the LOCAL group if the users are defined on the gateway itself. Enter the names of the AAA server groups; separate multiple entries with commas. You can click **Select** to select the groups or to create new ones.
Authentication Domain	A list or method for SSL VPN remote user authentication. If you do not specify a list or method, the gateway uses global AAA parameters for remote-user authentication.
Accounting Server Group	The accounting server group. Enter the name of the AAA server group policy object, or click **Select** to select it from a list or to create a new object.

SSL VPN Configuration Wizard—Portal Page Customization Page (IOS)

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Use this step of the SSL VPN Configuration wizard to define the appearance of the portal page that remote users see when connecting to the SSL VPN. The portal page allows remote users access to all websites available on the SSL VPN networks.

Navigation Path

(Device view) Open the Remote Access VPN Configuration Wizard for configuring a remote access SSL VPN on an ASA device (see Using the Remote Access VPN Configuration Wizard, page 30-13); then click **Next** until you reach this page.

Related Topics

- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32

Field Reference

Table 30-10 SSL VPN Configuration Wizard, Portal Page Customization Page

Element	Description
Title	The text displayed at the top of the page. Control the color using the Primary settings in the Title Color and Text Color fields.
Logo	The graphic displayed next to the title. Select None, Default, or Custom. To configure a custom graphic, you must copy the desired graphic to the Security Manager server, then click **Browse** to select the file. Supported graphic types are GIF, JPG, and PNG, with a maximum size of 100 KB.
Login Message	The text displayed immediately above the login prompt. Control the color using the Secondary settings in the Title Color and Text Color fields.
Title Color Text Color	The colors used for the title and login area and the text: • Primary—The Title, Login Box title, and the text in those areas. • Secondary—The Login Box username/password background and the text in that area. Click **Select** to choose background colors. For text, select either Black or White from the text list.
Preview	A preview of how the portal page will appear based on your selections.

Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices)

Note From version 4.17, though Cisco Security Manager continues to support Cisco IOS, FWSM, IPS, and PIX features/functionality, it does not support any enhancements.

This procedure describes how to create or edit IPSec VPNs on IOS and PIX 6.3 devices using the Remote Access VPN Configuration Wizard.

Tip The wizard allows you to select shared policies to use in the VPN on the Defaults page (the final step of the wizard). If you want to use this feature, you must first ensure that all required shared policies are configured and submitted to the database. For information on configuring shared policies and VPN policy defaults, see Understanding and Configuring VPN Default Policies, page 25-12.

Related Topics

- Understanding Remote Access IPSec VPNs, page 30-2
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8

Step 1 In Device view, select the desired IOS or PIX 6.3 device.

Step 2 From the Policy selector, select **Remote Access VPN > Configuration Wizard**.

Step 3 Select the **Remote Access IPSec VPN** radio button.

Step 4 Click **Remote Access Configuration Wizard**. The User Group Policy page opens.

Step 5 Select the required user groups from the Available User Groups list and click **>>**.

- If the required user group is not in the list, click **Create (+)** to open the Add User Groups dialog box, which enables you to create or edit a user group object. See Add or Edit User Group Dialog Box, page 34-73.
- You can edit an existing user group by selecting it in either list and clicking **Edit (pencil)**.
- To deselect a user group, select it and click **<<**.

Step 6 Click **Next**. The Defaults page opens.

Step 7 Select the shared policies to assign to the VPN. Initially, the selected policies are those configured on the Security Manager Administration VPN Defaults page. You can use the defaults or select different policies, if any are available. For more information about these policy defaults, see Remote Access VPN Configuration Wizard—Defaults Page, page 30-31.

Step 8 Click **Finish** to save your changes.

Inspect the policies created and configure any additional options that you want to implement.

Managing Remote Access VPNs on ASA and PIX 7.0+ Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

You can configure and manage remote access IPsec on devices running Cisco ASA Software or PIX 7.0+, and SSL VPNs on ASA 8.0+ devices (but not on PIX devices). Additionally, you can use IKE version 2 (IKEv2) negotiations in remote access IPsec VPNs on ASA 8.4(x) devices.

Tip No VPN configuration is supported on Cisco Catalyst 6500 Series ASA Service Modules and the ASA Software Release 8.5(x) used on the module.

The configuration of these remote access VPNs are the same for these device types. IOS and PIX 6.3+ devices use different configurations for remote access VPNs (as explained in Chapter 33, "Managing Remote Access VPNs on IOS and PIX 6.3 Devices").

The topics in this chapter explain how to configure policies that are specific to ASA and PIX 7.0+ devices. Additionally, review the following topics for more information about remote access VPNs:

- Understanding Remote Access VPNs, page 30-1
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8
- Discovering Remote Access VPN Policies, page 30-12
- Using the Remote Access VPN Configuration Wizard, page 30-13
  - Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (ASA and PIX 7.0+ Devices), page 30-25
  - Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (ASA Devices), page 30-14
- Chapter 32, "Managing Dynamic Access Policies for Remote Access VPNs (ASA 8.0+ Devices)"

This chapter contains the following topics:

- Overview of Remote Access VPN Policies for ASA and PIX 7.0+ Devices, page 31-2
- Understanding Cluster Load Balancing (ASA), page 31-5
- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Configuring Group Policies for Remote Access VPNs, page 31-26

- Understanding SSL VPN Server Verification (ASA), page 31-30
- Add/Edit Scripts Dialog Box, page 31-34
- Working with IPSec VPN Policies, page 31-36
- Working with SSL and IKEv2 IPSec VPN Policies, page 31-43
- Customizing Clientless SSL VPN Portals, page 31-77

Overview of Remote Access VPN Policies for ASA and PIX 7.0+ Devices

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

When you configure remote access VPNs on ASA or PIX 7.0+ devices, you use the following policies based on the type of VPN you are configuring. Possible remote access VPN types are: IKE version 1 (IKEv1) IPsec, IKE version 2 (IKEv2) IPsec, and SSL. IKEv2 is supported on ASA devices running the software version 8.4(x) and higher. Table 31-1 explains the conditions under which these policies are required or optional.

Note You cannot configure SSL VPNs on PIX devices; PIX devices support remote access IKEv1 IPsec VPNs only.

- **Policies used with remote access IKEv1 and IKEv2 IPsec and SSL VPNs:**
  - **ASA Cluster Load Balancing**—In a remote client configuration in which you are using two or more devices connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called load balancing. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. Load balancing is effective only on remote sessions initiated with an ASA device. For more information, see Understanding Cluster Load Balancing (ASA), page 31-5.
  - **Connection Profiles**—A connection profile is a set of records that contain VPN tunnel connection policies, including the attributes that pertain to creating the tunnel itself. Connection profiles identify the group policies for a specific connection, which includes user-oriented attributes. For more information, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.
  - **Dynamic Access**—Multiple variables can affect each VPN connection, for example, intranet configurations that frequently change, the various roles that each user might inhabit within an organization, and logins from remote access sites with different configurations and levels of security. Dynamic access policies (DAP) let you configure authorization that addresses these many variables. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. For more information, see Chapter 32, "Managing Dynamic Access Policies for Remote Access VPNs (ASA 8.0+ Devices)".

Note For multi-context ASA devices, the Dynamic Access policy is supported by Security Manager version 4.12 and ASA version 9.6(2) onwards only.

- **Global Settings**—You can define global settings that apply to all devices in your remote access VPNs. These settings include Internet Key Exchange (IKE), IKEv2, IPsec, NAT, and fragmentation definitions. The global settings typically have defaults that work in most situations, so configuring the Global Settings policy is optional in most cases; configure it only if you need non-default behavior or if you are supporting IKEv2 negotiations. For more information, see Configuring VPN Global Settings, page 26-30.
- **Group Policies**—You can view the user group policies defined for your remote access VPN connection profiles. From this page, you can specify new ASA user groups and edit existing ones. When you create a connection profile, if you specify a group policy that has not been used on the device, the group policy is automatically added to the Group Policies page; you do not need to add it to this policy before you create the connection profile. For more information, see Configuring Group Policies for Remote Access VPNs, page 31-26.
- **Public Key Infrastructure**—You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and to manage keys and certificates. Certification Authority (CA) servers are used to manage these certificate requests and issue certificates to users who connect to your IPsec or SSL remote access VPN. For more information, see Understanding Public Key Infrastructure Policies, page 26-51 and Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.

Note For multi-context ASA devices, the Public Key Infrastructure policy is supported by Security Manager version 4.12 and ASA version 9.6(2) onwards only.

- **Username from Cert Scripts**—You can use this policy to define a script to use in mapping the username from the certificate. For more information, see Add/Edit Scripts Dialog Box, page 31-34.

Note For multi-context ASA devices, the Username from Cert Scripts policy is supported by Security Manager version 4.12 and ASA version 9.6(2) onwards only.

- **Policies used in remote access IPsec VPNs only:**
  - **Certificate To Connection Profile Maps, Policy and Rules** (IKEv1 IPSec only.)—Certificate to connection profile map policies let you define rules to match a user's certificate to a permission group based on specified fields. To establish authentication, you can use any field of the certificate, or you can have all certificate users share a permission group. You can match the group from the DN rules, the Organization Unit (OU) field, the IKE identity, or the peer IP address. You can use any or all of these methods. For more information, see Configuring Certificate to Connection Profile Map Policies (ASA), page 31-36.
  - **IKE Proposal**—Internet Key Exchange (IKE), also called ISAKMP, is the negotiation protocol that enables two hosts to agree on how to build an IPsec security association. IKE is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs). Use the IKE Proposal policy to define the requirements for phase 1 of the IKE negotiation. For more information, see Configuring an IKE Proposal, page 26-9.
  - **IPsec Proposal (ASA/PIX 7.x)**—An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA. The policy is used for IKE phase 2 negotiations. For more information, see Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40.

- **Policies used in remote access IKEv2 IPSec and SSL VPNs only:**
  - **Access**—An Access policy specifies the security appliance interfaces on which a remote access SSL or IKEv2 IPsec VPN connection profile can be enabled, the port to be used for the connection profile, Datagram Transport Layer Security (DTLS) settings, the SSL VPN session timeout and maximum number of sessions. You can also specify whether to use the AnyConnect VPN Client or AnyConnect Essentials Client. For more information, see Understanding SSL VPN Access Policies (ASA), page 31-44.
  - **Other Settings**—The SSL VPN Other Settings policy defines settings that include caching, content rewriting, character encoding, proxy and proxy bypass definitions, browser plug-ins, AnyConnect client images and profiles, Kerberos Constrained Delegation, and some other advanced settings. For more information, see Configuring Other SSL VPN Settings (ASA), page 31-51.
  - **Shared License**—Use the SSL VPN Shared License page to configure your SSL VPN Shared License. For more information, see Configuring SSL VPN Shared Licenses (ASA 8.2+), page 31-74.

The following table explains whether a policy is required or optional for a particular type of VPN.

Table 31-1 Remote Access VPN Policy Requirements for ASA Devices

Policy	Required, Optional
ASA Cluster Load Balancing	**Optional** for all VPN types.
Connection Profiles	**Required** for all VPN types.
Dynamic Access	**Optional** for all VPN types.
Global Settings	**Required**: IKEv2 IPsec. **Optional**: IKEv1 IPsec, SSL.
Group Policies	**Required** for all VPN types.
Public Key Infrastructure	**Required**: IKEv2 IPsec. Also required if you configure any trustpoints for IKEv1 IPsec or SSL VPNs. Otherwise, it is optional.
Certificate To Connection Profile Maps, Policy and Rules	**Optional**: IKEv1 IPsec. **Not used in**: IKEv2 IPsec, SSL.
IKE Proposal	**Required**: IKEv1 IPsec, IKEv2 IPsec. **Not used in**: SSL.
IPsec Proposal (ASA/PIX 7.x)	**Required**: IKEv1 IPsec, IKEv2 IPsec. **Not used in**: SSL.
Access	**Required**: IKEv2 IPsec, SSL. **Not used in**: IKEv1 IPsec.
Other Settings	**Required**: IKEv2 IPsec, SSL. **Not used in**: IKEv1 IPsec.
Shared License	**Optional**: IKEv2 IPsec, SSL. **Not used in**: IKEv1 IPsec.

Understanding Cluster Load Balancing (ASA)

In a remote client configuration in which you are using two or more devices connected to the same network to handle remote sessions, you can configure these devices to share their session load. This feature is called load balancing. Load balancing directs session traffic to the least loaded device, thus distributing the load among all devices. Load balancing is effective only on remote sessions initiated with an ASA device.

To implement load balancing, you must group two or more devices on the same private LAN-to-LAN network into a virtual cluster. All devices in the virtual cluster carry session loads. One device in the virtual cluster, called the virtual director, directs incoming calls to the other devices, called secondary devices. The virtual director monitors all devices in the cluster, keeps track of how busy each is, and distributes the session load accordingly.

The virtual cluster appears to outside clients as a single virtual cluster IP address. This IP address is not tied to a specific physical device—it belongs to the current virtual director. A VPN client trying to establish a connection connects first to this virtual cluster IP address. The virtual director then sends back to the client the public IP address of the least-loaded available host in the cluster. In a second transaction (transparent to the user), the client connects directly to that host. In this way, the virtual director directs traffic evenly and efficiently across resources.

The role of virtual director is not tied to a physical device—it can shift among devices. If a machine in the cluster fails, the terminated sessions can immediately reconnect to the virtual cluster IP address. The virtual director then directs these connections to another active device in the cluster. Should the virtual director itself fail, a secondary device in the cluster immediately takes over as the new virtual session director. Even if several devices in the cluster fail, users can continue to connect to the cluster as long as any one device in the cluster is available.

Understanding Redirection Using a Fully Qualified Domain Name (FQDN)

By default, the ASA sends only IP addresses in load-balancing redirection to a client. If certificates are in use that are based on DNS names, the certificates will be invalid when redirected to a secondary device. As a VPN director, this security appliance can send a fully qualified domain name (FQDN) of a cluster device (another security appliance in the cluster) when redirecting VPN client connections to that cluster device. The security appliance uses reverse DNS lookup to resolve the FQDN of the device to its outside IP address to redirect connections and perform VPN load balancing. All outside and inside network interfaces on the load-balancing devices in a cluster must be on the same IP network.

After you enable load balancing using FQDNs, add an entry for each of your ASA outside interfaces into your DNS server, if such entries are not already present. Each ASA outside IP address should have a DNS entry associated with it for lookups. These DNS entries must also be enabled for Reverse Lookup. Enable DNS lookups on your ASA and define your DNS server IP address on the ASA.

For the procedure to configure cluster load balancing, see Configuring Cluster Load Balance Policies (ASA), page 31-5.

Configuring Cluster Load Balance Policies (ASA)

Use the ASA Cluster Load Balance page to enable load balancing for an ASA device in your remote access VPN. You must explicitly enable load balancing, as it is disabled by default. All devices that participate in a cluster must share the same cluster-specific values: IP address, encryption settings, encryption key, and port. For more information on cluster load balancing, see Understanding Cluster Load Balancing (ASA), page 31-5.

Note Load balancing requires an active 3DES/AES license and an ASA Model 5510 with a Plus license or an ASA Model 5520 or higher. The ASA device checks for the existence of this crypto license before enabling load balancing. If it does not detect an active 3DES or AES license, the device prevents load balancing, and also prevents internal configuration of 3DES by the load balancing system unless the license permits this usage.

Step 1 Do one of the following:

- (Device View) Select an ASA device; then select **Remote Access VPN > ASA Cluster Load Balance** from the Policy selector.
- (Policy View) Select **Remote Access VPN > ASA Cluster Load Balance** from the Policy Type selector. Select an existing policy or create a new one.

The ASA Cluster Load Balance page opens.

Step 2 Select **Participate in Load Balancing Cluster** to indicate that the device belongs to a load-balancing cluster.

Step 3 Configure the VPN Cluster Configuration options:

- **Cluster IPv4/IPv6 Address**—Specify the single IP address that represents the entire virtual cluster. Choose an IP address that is in the same subnet as the external interface. Beginning with version 4.12, Security Manager supports IPv6 address for IPv6 cluster in addition to the IPv4 address. This is for ASA devices running the version 9.0 or later.
- **UDP Port**—Specify the UDP destination port for the virtual cluster to which the device belongs. The port is typically 9023, but if that port is in use by another application, enter the UDP destination port number that you want to use for load balancing.
- **Enable IPSec Encryption**, **IPSec Shared Secret**—If required, select **Enable IPsec Encryption** to ensure that all load-balancing information communicated between the devices is encrypted. If you select this option, also enter (and confirm) the shared secret password. This can be a case-sensitive value between 4 and 16 characters, without spaces. The security appliances in the virtual cluster communicate through LAN-to-LAN tunnels using IPsec. This password must match the passwords passed on by the client.

Step 4 Configure the priority of the server in the cluster. Select one of the following options:

- **Accept default device value**—To accept the default priority value assigned to the device.
- **Configure same priority on all devices in the cluster**—To configure the same priority value to all the devices in the cluster. Then enter the priority number (1-10) to indicate the likelihood of the device becoming the virtual director, either at startup or when the existing director fails.

Step 5 Specify the public and private interfaces to be used on the server:

- **Public Interfaces**—The public interfaces to be used on the server. Enter the name of an interface or interface role object, or click **Select** to select the interface or role or to create a new role.
- **Private Interfaces**—The private interfaces to be used on the server. Enter the name of an interface or interface role object, or click **Select** to select the interface or role or to create a new role.

Step 6 If required, select **Send FQDN to client instead of an IP address when redirecting** to enable redirection using fully-qualified domain names. This option is available only for ASA devices running 8.0(2) or later. For more information, see Understanding Cluster Load Balancing (ASA), page 31-5.

Configuring Connection Profiles (ASA, PIX 7.0+)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

A connection profile is a set of records that contain VPN tunnel connection policies, including the attributes that pertain to creating the tunnel itself. Connection profiles identify the group policies for a specific connection, which includes user-oriented attributes. If you do not assign a group policy to a user, the default connection profile for the connection applies. You can create one or more connection profiles specific to your environment. You can configure connection profiles on the local remote access VPN server or on external AAA servers.

When you discover remote access VPN policies on a device, Security Manager adds the default connection profiles to the policy. You can edit these profiles, and the associated DlftGrpPolicy (renamed in Security Manager as <*device_display_name*>DfltGrpPolicy), but you cannot delete them. The following default connection profiles are supported in Security Manager:

- DefaultRAGroup—The default connection profile for remote access IPsec VPNs.
- DefaultWEBVPNGroup—The default connection profile for SSL VPNs. This connection profile is discovered only for ASA 8.0+ devices.

If you are configuring a connection profile on an ASA device, you have the option of configuring double authentication. The double authentication feature implements two-factor authentication for remote access to the network, in accordance with the Payment Card Industry Standards Council Data Security Standard. This feature requires that the user enter two separate sets of login credentials at the login page. For example, the primary authentication might be a one-time password, and the secondary authentication might be a domain (Active Directory) credential. If the primary credential authentication fails, the security appliance does not attempt to validate the secondary credentials. If either authentication fails, the connection is denied. Both the AnyConnect VPN client (SSL VPN or IKEv2 IPSec VPN) and Clientless SSL VPN support double authentication. The AnyConnect client supports double authentication on Windows computers (including supported Windows Mobile devices and Start Before Login), Mac computers, and Linux computers.

This procedure describes how to create or edit connection profiles on your remote access VPN server using the Connection Profile policy.

Note You can also create connection profiles from the Remote Access VPN Configuration wizard; see Using the Remote Access VPN Configuration Wizard, page 30-13. For information on connection profiles in Easy VPN site-to-site topologies, see Configuring a Connection Profile Policy for Easy VPN, page 28-13.

Related Topics

- Discovering Remote Access VPN Policies, page 30-12

Step 1 Do one of the following:

- (Device view) With an ASA or PIX 7.0+ device selected, select **Remote Access VPN > Connection Profiles** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Connection Profiles (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Connection Profiles page opens. The policy lists all connection profiles and shows the group policy used in the profile. For more information, see Connection Profiles Page, page 31-8.

Step 2 Click **Add Row (+)** beneath the table, or select a profile and click **Edit Row (pencil)**. The Connection Profiles dialog box opens.

Step 3 (All remote access VPN types.) On the General tab, specify the connection profile name and group policies and select which method (or methods) of address assignment to use. For a detailed explanation of the configuration, see General Tab (Connection Profiles), page 31-10.

Step 4 (All remote access VPN types.) Click the **AAA** tab to specify the AAA authentication parameters for the connection profile. For a detailed explanation of the configuration, see AAA Tab (Connection Profiles), page 31-13.

Step 5 (Remote access IKEv2 IPsec and SSL VPN only.) If you are setting up a connection profile on an ASA device, you can configure secondary authentication. To do so, click the **Secondary AAA** tab. For a detailed explanation of the configuration, see Secondary AAA Tab (Connection Profiles), page 31-17.

Step 6 (Remote access IPsec VPN only.) Click the **IPsec** tab to specify IPsec and IKE parameters for the connection profile. Some of these settings apply to IKEv1 but not to IKEv2 connections. For a detailed explanation of the configuration, see IPSec Tab (Connection Profiles), page 31-19.

Note To configure IKEv2 settings, use the IKEv2 Settings tab of the Global Settings policy; see Configuring VPN Global IKEv2 Settings, page 26-37.

Step 7 (Remote access SSL VPN only.) Click the **SSL** tab to specify the WINS servers for the connection profile policy, select a customized look and feel for the SSL VPN end-user logon web page, specify DHCP servers to be used for client address assignment, and establish an association between an interface and client IP address pools. For a detailed explanation of the configuration, see SSL Tab (Connection Profiles), page 31-22.

Step 8 Click **OK**.

Connection Profiles Page

Use the Connection Profiles page to manage connection profile policies for remote access VPN or Easy VPN topologies. The Connection Profiles page lists the connection profiles that are configured, shows the group policy associated with those connection profiles, and indicates whether a connection profile is the default connection profile to use for Citrix clients when no specific tunnel group is identified during tunnel negotiation.

Use of this policy differs depending on the type of VPN you are configuring:

- Remote access SSL VPN—The policy is used only for ASA devices. You can create multiple profiles, and configure settings on all tabs of the Connection Profiles dialog box.
- Remote access IPSec VPN—The policy is used for ASA devices and PIX Firewalls running PIX 7.0+ software. You can create multiple profiles, but only the General, AAA, and IPSec tabs on the Connection Profiles dialog box apply to this configuration (in some cases, you will see only these tabs).
- Easy VPN topologies—The policy is used for Easy VPN servers (hubs) that are ASA devices or PIX Firewalls running PIX 7.0+ software. You can create a single profile, so the policy page actually imbeds the Connection Profiles dialog box, so that you have direct access to the tabs that define the profile. Only the General, AAA, and IPSec tabs apply.

For remote access IPSec and SSL VPNs:

- To add a profile, click the **Add Row** button and fill in the Connection Profiles dialog box.
- To edit an existing profile, select it and click the **Edit Row** button.
- To delete a profile, select it and click the **Delete Row** button.

The connection profile consists of the following tabs. Configure them as appropriate for the type of VPN you are configuring.

- General Tab (Connection Profiles), page 31-10
- AAA Tab (Connection Profiles), page 31-13
- Secondary AAA Tab (Connection Profiles), page 31-17 (SSL VPN and IKEv2 IPsec VPN only.)
- IPSec Tab (Connection Profiles), page 31-19 (Some of these settings apply to IKEv1 but not to IKEv2 connections.)
- SSL Tab (Connection Profiles), page 31-22 (SSL VPN only)

Navigation Path

Remote access VPNs:

- (Device View) Select an ASA or PIX 7+ device and select **Remote Access VPN > Connection Profiles** from the Policy selector.
- (Policy View) Select **Remote Access VPN > Connection Profiles (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Easy VPN:

- From the Site-to-Site VPN Manager Window, page 25-18, select the Easy VPN topology and then select **Connection Profiles (PIX7.0/ASA)**.
- (Device view) Select a device that participates in the Easy VPN topology and select **Site to Site VPN** from the Policy selector. Select the Easy VPN topology and click **Edit VPN Policies** to open the Site-to-Site VPN Manager Window, page 25-18, where you can select the policy.
- (Policy view) Select **Site-to-Site VPN > Connection Profiles (PIX7.0/ASA)**. Select an existing policy or create a new one.

This section contains the following topics:

- General Tab (Connection Profiles), page 31-10
- AAA Tab (Connection Profiles), page 31-13
- Secondary AAA Tab (Connection Profiles), page 31-17
- IPSec Tab (Connection Profiles), page 31-19
- SSL Tab (Connection Profiles), page 31-22

Supported CLIs in Remote Access VPN Multi-Context Mode - Connection Profiles

The following CLIs are supported for ASA 9.5(2) for Connection Profiles for remote access VPN in multiple context mode. These CLIs are supported in Admin and User Context for Tunnel-Group.

DefaultWEBVPNGroup is the default Connection Profile. DefaultRAGroup is not supported in ASA 9.5(2) remote access VPN Multiple Context mode.

Note For the configurations that are not supported, Security Manager displays a warning message that you can ignore. No delta will be generated.

- Type remote-access
- General-attributes
  - Accounting-server-group
  - Address-pool
  - Annotation
  - Authenticated-session-username
  - Authentication-attr-from-server
  - Authentication-server-group
  - Authorization-required
  - Authorization-server-group
  - Default-group-policy
  - Dhcp-server
  - Exit
  - Ipv6-address-pool
  - Nat-assigned-to-public-ip
  - Password-management
  - Secondary-authentication-server-group
- Webvpn-attributes
  - Authentication
  - Exit
  - Group-alias
  - Group-url
  - No
  - Radius-reject-message

General Tab (Connection Profiles)

Use the General tab of the Connection Profiles dialog box to configure the basic properties for a VPN Connection Profile policy. These properties are used in remote access IPsec and SSL VPNs and site-to-site Easy VPN topologies.

The General Tab is supported in ASA 9.5(2) Remote Access VPN in Multiple Context mode.

Navigation Path

- Remote Access VPNs—From the Connection Profiles page (see Connection Profiles Page, page 31-8), click the **Add Row (+)** button, or select a profile and click the **Edit Row (pencil)** button, to open the Connection Profiles dialog box. Click the **General** tab if necessary.

- Easy VPN topologies—Select the site-to-site VPN Connection Profiles policy in either Policy view or in the Site-to-Site VPN Manager with an Easy VPN topology selected (see Connection Profiles Page, page 31-8). Click the **General** tab if necessary.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- ASA Group Policies Dialog Box, page 34-1
- Understanding Networks/Hosts Objects, page 6-80
- Configuring a Connection Profile Policy for Easy VPN, page 28-13
- Understanding Easy VPN, page 28-1

Field Reference

Table 31-2 Connection Profile General Tab

Element	Description
Connection Profile Name	The name of the connection profile (tunnel group).
Group Policy	If required, the name of the ASA group policy object that defines the default user group associated with the connection profile. A group policy is a collection of user-oriented attribute/value pairs stored either internally on the device or externally on a RADIUS/LDAP server. Click **Select** to select an existing object or to create a new one.
Client Address Assignment	
DHCP Servers	The DHCP servers to be used for client address assignments. The servers are used in the order listed. Enter the IP addresses of the DHCP servers or the names of network/host policy objects that define the DHCP server addresses. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.
Global IPv4 Address Pool	The address pools from which IPv4 addresses will be assigned to clients if no pool is specified for the interface to which the client connects. Address pools are entered as a range of addresses, such as 10.100.12.2-10.100.12.254. The server uses these pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools. Enter the address pool ranges or the names of network/host objects that define these pools. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.

Table 31-2 Connection Profile General Tab (continued)

Element	Description
Global IPv6 Address Pool	The address pools from which IPv6 addresses will be assigned to clients if no pool is specified for the interface to which the client connects. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running version 9.0 or later. Address pools are entered as a range of addresses, for example, fe80::60/5 4, where fe80::60/5 is the IPv6 address and prefix length, and 4 is the count (number of addresses). The server uses these pools in the order listed. If all addresses in the first pool have been assigned, it uses the next pool, and so on. You can specify up to 6 pools. Enter the address pool ranges or the names of network/host objects that define these pools. Click **Select** to select existing network/host objects or to create new ones. Separate multiple entries with commas.
Interface-Specific Address Pools table	If you want to configure separate IP address pools for specific interfaces, so that clients connecting through that interface use a pool different from the global pool, add the interface to this table and configure the separate pool. Any interface not listed here uses the global pool. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running version 9.0 or later. Therefore, you see an additional column for IPv6 address pool. • To add an interface-specific address pool, click the **Add Row** button and fill in the Add/Edit Interface Specific Client Address Pools Dialog Box, page 31-12 • To edit an interface pool, select it and click the **Edit Row** button. • To delete an interface, select it and click the **Delete Row** button.

Add/Edit Interface Specific Client Address Pools Dialog Box

Use the Add/Edit Interface Specific Client Address Pools dialog box to configure interface-specific client address pools for your connection profile policy.

Navigation Path

Open the General tab in the Connection Profiles dialog box (see General Tab (Connection Profiles), page 31-10), then click **Add Row** below the Interface-Specific Address Pools table, or select a row in the table and click **Edit Row**.

Related Topics

- Creating Interface Role Objects, page 6-74
- Creating Networks/Hosts Objects, page 6-82

Field Reference

Table 31-3 Add/Edit Interface Specific Client Address Pools Dialog Box

Element	Description
Interface	The interface to which you are assigning an address pool. Enter the interface name or the name of an interface role object, or click **Select** to select an interface or object or to create a new object.
IPv4 Address Pool	The IPv4 address pool to assign to the interface. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running the version 9.0 or later. Address pools are specified using the starting and ending IP addresses of the pool, for example, 10.100.10.2-10.100.10.254. You can either type in the IP address range, or use a network/host object that specifies an address range. Click **Select** to select a network/host object or to create a new object.
IPv6 Address Pool	TheIPv6 address pool to assign to the interface. IPv6 address pools are specified using the IPv6 address along with the prefix length and followed by the count, where count refer to the number of addresses in the pool. You can either type in the IP address range, or use a network/host object that specifies an address range. Click **Select** to select a network/host object or to create a new object.

AAA Tab (Connection Profiles)

Use the AAA tab of the Connection Profile dialog box to configure the AAA authentication parameters for a connection profile policy.

For AAA, the Distinguished Name Authorization Settings policy is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

However, beginning with Security Manager version 4.12, this policy is supported for ASA 9.6(2) Remote Access VPN in Multi-context mode. The supported CLIs for the Admin and User context are:

- Tunnel-group General-attributes
  - Secondary-username-from-certificate
  - Username-from-certificate

Navigation Path

- Remote Access VPNs—From the Connection Profiles page (see Connection Profiles Page, page 31-8), click the **Add Row (+)** button, or select a profile and click the **Edit Row (pencil)** button, to open the Connection Profiles dialog box. Click the **AAA** tab.
- Easy VPN topologies—Select the site-to-site VPN Connection Profiles policy in either Policy view or in the Site-to-Site VPN Manager with an Easy VPN topology selected (see Connection Profiles Page, page 31-8). Click the **AAA** tab.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Understanding AAA Server and Server Group Objects, page 6-27
- Configuring a Connection Profile Policy for Easy VPN, page 28-13
- Understanding Easy VPN, page 28-1

Field Reference

Table 31-4 Connection Profile AAA Tab

Element	Description
Authentication Method	Whether to authenticate connections using AAA, certificates, or both, Multiple Certificate, AAA and Multiple Certificate, and SAML. If you select Certificate, many of the options on the dialog box are disabled as the required details are obtained from the certificate. Beginning with version 4.10, Security Manager enables you to select SAML Identity Provider as an authentication method. This is to enable SAML Service Provider for the current tunnel group. SAML Identity Provider will not be used until they are applied in a tunnel group. The SAML authentication is a mutual exclusion authentication method. See Configuring SAML Identity Provider, page 6-98 for more information. Beginning with version 4.13, Security Manager enables you to select Multiple Certificate, or AAA & Multiple Certificate as an authentication method. This method is enabled to support the Multiple certificate authentication feature of ASA 9.7.1 devices. If you select this method for a ASA device earlier to 9.7.1 release, validation error message appears. For more information, see Multiple Certificate Authentication Support, page 9-4.
Authentication Server Group	The name of the authentication server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object. If you want to use different authentication server groups based on the interface to which the client connects, configure the server groups in the Interface-Specific Authentication Server Groups table at the bottom of this tab (described below).
Use LOCAL if Server Group Fails	Whether to fall back to the local database for authentication if the selected authentication server group fails.
Authorization Server Group	The name of the authorization server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Users must exist in the authorization database to connect	Whether to require that the username of the client must exist in the authorization database to allow a successful connection. If the username does not exist in the authorization database, then the connection is denied.
Accounting Server Group	The name of the accounting server group. Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object.
Strip Realm from Username Strip Group from Username	Whether to remove the realm or group name from the username before passing the username on to the AAA server. A realm is an administrative domain. Enabling these options allows the authentication to be based on the username alone. You can enable any combination of these options. However, you must select both check boxes if your server cannot parse delimiters.

Table 31-4 Connection Profile AAA Tab (continued)

Element	Description
Override Account-Disabled Indication from AAA Server	Whether to override the "account-disabled" indicator from a AAA server. This configuration is valid for servers, such as RADIUS with NT LDAP, and Kerberos, that return an "account-disabled" indication. If you are using an LDAP directory server for authentication, password management is supported with the Sun Microsystems JAVA System Directory Server (formerly named the Sun ONE Directory Server) and the Microsoft Active Directory. • **Sun**—The DN configured on the security appliance to access a Sun directory server must be able to access the default password policy on that server. We recommend using the directory administrator, or a user with directory administrator privileges, as the DN. Alternatively, you can place an ACI on the default password policy. • **Microsoft**—You must configure LDAP over SSL to enable password management with Microsoft Active Directory.
Enable Notification Upon Password Expiration to Allow User to Change Password Enable Notification Prior to Expiration Notify Prior to Expiration	Whether to have the security appliance notify the remote user at login that the current password is about to expire or has expired, and to then offer the user the opportunity to change the password. If you want to give the user prior warning of an impending password expiration, select **Enable Notification Prior to Expiration** and specify the number of days prior to expiration that you want to start notifications (1 to 180 days). You can use this option with AAA servers that support such notification—RADIUS, RADIUS with an NT server, and LDAP servers. There is no prior notification for other types of servers.
Distinguished Name (DN) Authorization Settings	How you want to use the distinguished name for authorization. A distinguished name (DN) is a unique identification, made up of individual fields, that can be used as the identifier when matching users to a tunnel group. DN rules are used for enhanced certificate authentication. Select from the following options to determine how the DN is used during authorization: • **Use Entire DN as the Username**—Use the entire DN; do not focus on any one field. • **Specify Individual DN fields as the Username**—Focus on specific fields. Select a primary field, and optionally, a secondary field. The default is to use the common name (CN) as primary and the organization unit (OU) as secondary. • **Use Script to Select Username**—Beginning with version 4.7, Security Manager enables you to define a script to use in mapping the username from the certificate. Select the script that you have defined from the drop-down list. For more information, see Add/Edit Scripts Dialog Box, page 31-34.

Table 31-4 Connection Profile AAA Tab (continued)

Element	Description
Interface-Specific Authentication Server Groups table	If you want to configure separate authentication server groups for specific interfaces, so that clients connecting through that interface use a server group different from the global group, add the interface to this table and configure the separate group. Any interface not listed here uses the global authentication server group. The table shows the server group and whether you are falling back to local authentication if the server group is not available. • To add an interface-specific authentication group to the list, click the **Add Row** button and fill in the Add/Edit Interface Specific Authentication Server Groups Dialog Box, page 31-16. • To edit an interface setting, select it and click the **Edit Row** button. • To delete an interface setting, select it and click the **Delete Row** button.

Add/Edit Interface Specific Authentication Server Groups Dialog Box

Use the Add/Edit Interface Specific Authentication Server Groups dialog boxes to configure interface-specific authentication for your connection profile policy. This setting overrides the global authentication server group settings if the client connects to the specified interface.

If you are configuring the secondary AAA server for an SSL VPN on an ASA device, the settings are specifically used for the secondary set of credentials that the user enters; this is reflected in the name of the dialog box.

Navigation Path

Open the AAA or Secondary AAA tabs in the Connection Profiles dialog box (see AAA Tab (Connection Profiles), page 31-13 or Secondary AAA Tab (Connection Profiles), page 31-17), then click **Add Row** below the Interface-Specific Address Pools table, or select a row in the table and click **Edit Row**.

Related Topics

- Understanding Interface Role Objects, page 6-73
- Understanding AAA Server and Server Group Objects, page 6-27

Field Reference

Table 31-5 Add/Edit (Secondary) Interface Specific Authentication Server Groups

Element	Description
Interface	The name of the interface or interface role (that identifies the interfaces) for which you are configuring an authentication server group. Click **Select** to select an interface or interface role or to create a new interface role.

Table 31-5 Add/Edit (Secondary) Interface Specific Authentication Server Groups (continued)

Element	Description
Server Group	The name of the authentication server group (LOCAL if the tunnel group is configured on the local device). Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object. When you are configuring secondary AAA, this group is used specifically for the second credentials. You can specify different server groups for primary and secondary credentials.
Use LOCAL if Server Group Fails	Whether to fall back to the local database for authentication if the selected authentication server group fails.
Use Primary Username (Secondary authentication only; remote access SSL or IKEv2 IPSec VPN on ASA 8.2+ only.)	Whether to use the same username for the secondary credentials that was used for the primary credentials. If you select this option, after users authenticate with their primary credentials, they are prompted for the secondary password only. If you do not select this option, the secondary prompt requires both a username and password.

Secondary AAA Tab (Connection Profiles)

Use the Secondary AAA tab to configure the secondary AAA authentication parameters for a remote access SSL VPN connection profile policy for use with ASA 8.2+ devices, or a remote access IKEv2 IPSec VPN connection profile policy for use with an ASA 8.4(1)+ device. These settings do not apply to remote access IKEv1 IPSec VPNs or Easy VPN topologies or to other device types.

Navigation Path

Remote Access VPNs only—From the Connection Profiles page (see Connection Profiles Page, page 31-8), click the **Add Row (+)** button, or select a profile and click the **Edit Row (pencil)** button, to open the Connection Profiles dialog box. Click the **Secondary AAA** tab.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7

Field Reference

Table 31-6 Connection Profile Secondary AAA Tab

Element	Description
Enable Double Authentication	Whether to enable double authentication, which prompts the user for two sets of credentials (username and password) before completing the remote access VPN connection.
Secondary Authentication Server Group	The name of the authentication server group (LOCAL if the tunnel group is configured on the local device) to be used with the second set of credentials. Enter the name of a AAA server group object or click **Select** to select it from a list or to create a new object. If you want to use different authentication server groups based on the interface to which the client connects, configure the server groups in the Secondary Interface-Specific Authentication Server Groups table at the bottom of this tab (described below).

Table 31-6 Connection Profile Secondary AAA Tab (continued)

Element	Description
Use LOCAL if Server Group Fails	Whether to fall back to the local database for authentication if the selected authentication server group fails.
Use Primary Username for Secondary Authentication	Whether to use the same username for the secondary credentials that was used for the primary credentials. If you select this option, after users authenticate with their primary credentials, they are prompted for the secondary password only. If you do not select this option, the secondary prompt requires both a username and password.
Username for Session	The username that the software will use for the user session, either the primary or secondary name. If you prompt for the primary name only, select primary. **Note** By default, if there is more than one username, AnyConnect remembers both usernames between sessions. In addition, the head-end device might offer a feature to allow for administrative control over whether the client remembers both or neither usernames.
Authorization Authentication Server	The server to use for authorization, either the primary authentication server (defined on the AAA tab) or the secondary authentication server configured on this tab.

Distinguished Name (DN) Secondary Authorization Setting	How you want to use the distinguished name for authorization. A distinguished name (DN) is a unique identification, made up of individual fields, that can be used as the identifier when matching users to a tunnel group. DN rules are used for enhanced certificate authentication. Select from the following options to determine how the DN is used during authorization:
	• **Use Entire DN as the Username**—Use the entire DN; do not focus on any one field.
	• **Specify Individual DN fields as the Username**—Focus on specific fields. Select a primary field, and optionally, a secondary field. The default is to use only the user identification (UID) field.
	• **Use Script to Select Username**—Beginning with version 4.7, Security Manager enables you to define a script to use in mapping the username from the certificate. Select the script that you have defined from the drop-down list. For more information, see Add/Edit Scripts Dialog Box, page 31-34.
	Note The Distinguished Name (DN) Secondary Authorization Settings policy is supported from Security Manager version 4.12 for ASA devices running version 9.6(2) in Multi-context mode. The supported CLIs for the Admin and User context are:
	• Tunnel-group General-attributes
	  • Secondary-username-from-certificate
	  • Username-from-certificate

Table 31-6 Connection Profile Secondary AAA Tab (continued)

Element	Description
Secondary Interface-Specific Authentication Server Groups table	If you want to configure separate secondary authentication server groups for specific interfaces, so that clients connecting through that interface use a server group different from the global group, add the interface to this table and configure the separate group. Any interface not listed here uses the global authentication server group. The table shows the server group and whether you are falling back to local authentication if the server group is not available. • To add a secondary interface-specific authentication group to the list, click the **Add Row** button and fill in the Add/Edit Interface Specific Authentication Server Groups Dialog Box, page 31-16. • To edit an interface setting, select it and click the **Edit Row** button. • To delete an interface setting, select it and click the **Delete Row** button.

IPSec Tab (Connection Profiles)

Use the IPsec tab of the Connection Profiles page to specify IPsec and IKE parameters for the connection policy.

Beginning with version 4.8, Security Manager supports VPN connectivity via standards-based, third-party, IKEv2 remote-access clients (in addition to AnyConnect). Authentication support includes preshared keys, certificates, and user authentication via the Extensible Authentication Protocol (EAP).

IPSec is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode. Beginning with Cisco Security Manager version 4.17, IPSec is supported from ASA 9.9(2) or later multi-context devices. However, the following attributes under **Connection Profile > IPSec** tab is not supported for ASA 9.9(2) or later multi-context devices:

- Enable IKEv2 Mobike RRC
- Client Software Update Table

Navigation Path

- Remote Access VPNs—From the Connection Profiles page (see Connection Profiles Page, page 31-8), click the **Add Row (+)** button, or select a profile and click the **Edit Row (pencil)** button, to open the Connection Profiles dialog box. Click the **IPSec** tab.
- Easy VPN topologies—Select the site-to-site VPN Connection Profiles policy in either Policy view or in the Site-to-Site VPN Manager with an Easy VPN topology selected (see Connection Profiles Page, page 31-8). Click the **IPSec** tab.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Configuring a Connection Profile Policy for Easy VPN, page 28-13
- Understanding Easy VPN, page 28-1

Field Reference

Table 31-7 Connection Profiles IPsec Tab

Element	Description
IKEv1 Peer Authentication	
Preshared Key	The preshared key for the connection profile. The maximum length of a preshared key is 127 characters. Enter the key again in the Confirm field.
Trustpoint Name	The name of the PKI enrollment policy object that defines the trustpoint name if any trustpoints are configured for IKEv1 connections. A trustpoint represents a Certificate Authority (CA)/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate. Click **Select** to select the object from a list or to create a new object. **Tips** If you specify a trustpoint, you must also select the same PKI enrollment object in the Public Key Infrastructure policy. For more information, see Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.
IKEv2 Peer Authentication You can configure one or more authentication options such as preshared key, certificate, and EAP for remote authentication.	
Preshared Key	The preshared key for the connection profile. The maximum length of a preshared key is 127 characters. Enter the key again in the Confirm field.
Enable Certificate Authentication	Allows you to use certificates for authentication if checked.
Enable EAP Authentication	Allows you to use EAP for authentication if checked. **Note** You must use certificates for local authentication if you check this check box since EAP authentication requires the server to authenticate via a certificate.
Send EAP identity request to the client	Enables you to send an EAP request for authentication to the remote access VPN client.
IKEv2 Local Authentication You can configure either a preshared key or a trustpoint name for local authentication.	
Preshared key	The preshared key for the connection profile. The maximum length of a preshared key is 127 characters. Enter the key again in the Confirm field.

Table 31-7 Connection Profiles IPsec Tab (continued)

Element	Description
Trustpoint Name	The name of the PKI enrollment policy object that defines the trustpoint name if any trustpoints are configured for IKEv2 connections. A trustpoint represents a Certificate Authority (CA)/identity pair and contains the identity of the CA, CA-specific configuration parameters, and an association with one enrolled identity certificate. Click **Select** to select the object from a list or to create a new object. **Note** Local authentication must be using certificate if you select EAP for remote authentication.
Enable IKEv2 Mobike RRC	Select to enable Return Routability checking for dynamic IP address changes in IKE/IPSEC security associations on which mobike is enabled. By default, Mobike RRC is disabled. **Note** You can enable Return Routability Checking for dynamic IP address changes, only for ASA 9.8.1 and higher. **Note** This option is not supported for ASA 9.9(2) or later multi-context devices.
IKE Peer ID Validation	Select whether IKE peer ID validation is ignored (Do not check), required, or checked only if supported by a certificate. During IKE negotiations, peers must identify themselves to one another.
Enable Sending Certificate Chain	Whether to enable the sending of the certificate chain for authorization. A certificate chain includes the root CA certificate, identity certificate, and key pair.
Enable Password Update with RADIUS Authentication	Whether to enable passwords to be updated with the RADIUS authentication protocol. For more information, see Supported AAA Server Types, page 6-28.
ISAKMP Keepalive	Whether to monitor ISAKMP keepalive. If you select the **Monitor Keepalive** option, you can configure IKE keepalive as the default failover and routing mechanism. Enter the following parameters: • **Confidence Interval**—The number of seconds that a device waits between sending IKE keepalive packets. • **Retry Interval**—The number of seconds a device waits between attempts to establish an IKE connection with the remote peer. The default is 2 seconds. For more information, see Configuring VPN Global ISAKMP/IPsec Settings, page 26-33.
Client Software Update table	The VPN client revision level and URLs for client platforms. You can configure different revision levels for All Windows Platforms, Windows 95/98/ME, Windows NT4.0/2000/XP, or the VPN3002 Hardware Client. To configure the client for a platform, select it, click the **Edit Row** button, and fill in the IPSec Client Software Update Dialog Box, page 31-22. **Note** This option is not supported for ASA 9.9(2) or later multi-context devices.

IPSec Client Software Update Dialog Box

Use the IPsec Client Software Update dialog box to configure the specific revision level and image URL of a VPN client.

Navigation Path

Open the IPSec tab in the Connection Profiles dialog box (see IPSec Tab (Connection Profiles), page 31-19), select a client type from the Client Software Update table, then click **Edit Row**.

Field Reference

Table 31-8 IPSec Client Software Update Dialog Box

Element	Description
Client Type	Type of client being modified.
Client Revisions	Revision level of the client.
Image URL	URL of the client software image.

SSL Tab (Connection Profiles)

Use the SSL tab of the Connection Profile dialog box to configure the WINS servers for the connection profile policy, select a customized look and feel for the SSL VPN end-user logon web page, DHCP servers to be used for client address assignment, and to establish an association between an interface and client IP address pools. Some items, such as connection profile aliases, apply to remote access IKEv2 IPsec VPNs, but otherwise these settings do not apply to remote access IKEv1 IPSec VPNs or Easy VPN topologies.

The following profiles are supported for the SSL tab in ASA 9.5(2) Remote Access VPN in Multi-context mode.

- Radius-Reject-Message
- Connection alias
- Group-url
- Group-alias

Navigation Path

Remote Access VPNs only—From the Connection Profiles page (see Connection Profiles Page, page 31-8), click the **Add Row (+)** button, or select a profile and click the **Edit Row (pencil)** button, to open the Connection Profiles dialog box. Click the **SSL** tab.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs, page 31-88
- Understanding Networks/Hosts Objects, page 6-80
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 31-9 Connection Profile SSL Tab

Element	Description
WINS Servers List	The name of the WINS (Windows Internet Naming Server) servers list to use for CIFS name resolution. Click **Select** to select the WINS servers list policy object or to create a new object. SSL VPN uses the CIFS protocol to access or share files on remote systems. When you attempt a file-sharing connection to a Windows computer by using its computer name, the file server you specify corresponds to a specific WINS server name that identifies a resource on the network. A WINS servers list defines a list of WINS servers, which are used to translate Windows file server names to IP addresses. The security appliance queries the WINS servers to map WINS names to IP addresses. You must configure at least one, and up to three WINS servers for redundancy. The security appliance uses the first server on the list for WINS/CIFS name resolution. If the query fails, it uses the next server.
DNS Group	The DNS group to use for the SSL VPN tunnel group. The DNS group resolves the hostname to the appropriate DNS server for the tunnel group. Select the desired group from the list; the DefaultDNS group is the default group that is always available on the device. **Tip** The DNS groups are defined in the **Platform > Device Admin > Server Access > DNS** policy. Use the DNS policy to change the servers defined in a group or to add or remove groups. See DNS Page, page 52-14.
Portal Page Customization	The name of the SSL VPN Customization policy object that defines the default portal page for the VPN. This profile defines the appearance of the portal page that allows the remote user access to all resources available on the SSL VPN. Click **Select** to select the object or to create a new object. **Note** You can set up different login windows for different groups by using a combination of customization profiles and groups. For example, assuming that you had created a customization profile called salesgui, you can create an SSL VPN group called sales that uses that customization profile. You would then specify the SSL VPN customization object in the group policy object on the SSL VPN > Settings tab; see ASA Group Policies SSL VPN Settings, page 34-25.
SAML Identity Provider	Select the SAML Identity Provider. SAML Identity Provider will not be used until it is applied in a tunnel group. See Configuring SAML Identity Provider, page 6-98 for more information.
Override SVC Download (ASA 8.0(2)+ only)	Whether you want clientless users logging in under specific tunnel groups to not have to wait for the download prompt to expire before being presented with the clientless SSL VPN home page. Instead, these users are immediately presented with the clientless SSL VPN home page.

Table 31-9 Connection Profile SSL Tab (continued)

Element	Description
Reject Radius Message (ASA 8.0(2)+ only)	Whether you want to display to remote users a RADIUS message about their authentication failure.
Connection Aliases table	A list of alternate names by which the tunnel group can be referred to. The status indicates whether the name is enabled for use or disabled (cannot be used). A group alias creates one or more alternate names by which a user can refer to a tunnel group. This feature is useful when the same group is known by several common names (such as "Devtest" and "QA"). If you want the actual name of the tunnel group to appear on this list, you must specify it as an alias. The group alias that you specify here appears on the login page. Each tunnel group can have multiple aliases or no alias. • To add an alias, click the **Add Row (+)** button beneath the table and fill in the Add/Edit Connection Alias Dialog Box, page 31-25. • To edit an alias, select it and click the **Edit Row (pencil)** button. • To delete an alias, select it and click the **Delete Row (trash can)** button.
Group URLs table	A list of URLs associated with the tunnel group connection profile. The status indicates whether the URL is enabled for use. When enabled, the user can use the URL, which eliminates the need to select a group during login. You can configure multiple URLs (or no URLs) for a tunnel group. Each URL can be enabled or disabled individually. You must use a separate specification for each URL, specifying the entire URL using either the HTTP or HTTPS protocol. • To add a URL, click the **Add Row (+)** button beneath the table and fill in the Add/Edit Connection URL Dialog Box, page 31-25. • To edit a URL, select it and click the **Edit Row (pencil)** button. • To delete a URL, select it and click the **Delete Row (trash can)** button.
Default Citrix Client Profile (ASA 9.1(4)+ only)	Whether this connection profile should be the default connection profile to use for Citrix clients when no specific tunnel group is identified during tunnel negotiation. **Note** Only one connection profile can be configured as the Default Citrix Client Profile. If you try to configure a connection profile as the Default Citrix Client Profile when another profile is already configured as such, you will receive a warning message. If you continue with the operation, the selected connection profile will be made the Default Citrix Client Profile and the other connection profile will be deselected as the Default Citrix Client Profile.

Table 31-9 Connection Profile SSL Tab (continued)

Element	Description
Disable CSD (ASA 8.2(0)+ only) Both Clientless and AnyConnect AnyConnect only	Whether to disable Cisco Secure Desktop (CSD) for this connection profile. Security Manager supports this feature on all devices that are running ASA software version 8.2(0) and higher. **Note** If you choose to disable CSD, by default Security Manager selects the option for both SSL Clientless VPN and AnyConnect.

Add/Edit Connection Alias Dialog Box

Use the Add/Edit Connection Alias dialog box to create or edit a connection alias for an SSL or IKEv2 IPsec VPN connection profile. Specifying the connection alias creates one or more alternate names by which the user can refer to a tunnel group.

Navigation Path

Open the SSL tab in the Connection Profiles dialog box (see SSL Tab (Connection Profiles), page 31-22), and click **Add Row** beneath the Connection Alias table, or select an alias from the table and click **Edit Row**.

Field Reference

Table 31-10 Add/Edit Connection Alias Dialog Box

Element	Description
Enabled	Whether to enable the connection alias. You must enable the alias for users to use it.
Connection Alias	The alternative name for the connection profile. The connection alias that you specify here appears in a list on the user's login page.

Add/Edit Connection URL Dialog Box

Use this dialog box to specify incoming URLs for the tunnel group. If a connection URL is enabled in a tunnel group, when the user connects using that URL, the security appliance selects the associated tunnel group and presents the user with only the username and password fields in the login window.

Tips

- You can configure multiple URLs or addresses (or none) for a group. Each URL or address can be enabled or disabled individually.
- You cannot associate the same URL or address with multiple groups. The security appliance verifies the uniqueness of the URL or address before accepting the URL or address for a tunnel group.

Navigation Path

Open the SSL tab in the Connection Profiles dialog box (see SSL Tab (Connection Profiles), page 31-22), and click **Add Row** beneath the Group URLs table, or select a URL from the table and click **Edit Row**.

Field Reference

Table 31-11 Add/Edit Connection URL Dialog Box

Element	Description
Enabled	Whether to enable the connection alias. You must enable the alias for users to use it.
Connection URL	Select a protocol (**http** or **https**) from the list, and specify the incoming URL for the connection.

Configuring Group Policies for Remote Access VPNs

In the Group Policies page, you can view the user group policies defined for your ASA remote access VPN connection profiles. From this page, you can specify new ASA user groups and edit existing ones. When you create a connection profile, if you specify a group policy that has not been used on the device, the group policy is automatically added to the Group Policies page; you do not need to add it to this policy before you create the connection profile. For information on creating connection profiles, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.

For more information about group policies, see Understanding Group Policies (ASA), page 31-27.

Tip Dynamic Access policies take precedence over Group policies. If a setting is not specified in a Dynamic Access policy, an ASA device checks for Group policies that specify the setting.

Each row in the table represents an ASA group policy object, displaying the name of the policy object assigned to the remote access VPN connection profile, whether it is stored on the ASA device itself (Internal) or on a AAA server (External), and whether the group is for IKEv1 (IPsec), IKEv2 (IPsec), SSL, or all types of VPN. For external groups, the protocol is unknown and listed as N/A.

- To add an ASA group policy object, click the **Add Row** button. This opens an object selector, from which you can select an existing policy object or click the **Create** button to create a new object. For more information about creating group policies, see Creating Group Policies (ASA, PIX 7.0+), page 31-28.

  Note You cannot create more than one group policy that includes DfltGrpPolicy in its name. DfltGrpPolicy is the default policy defined on the device; if Security Manager discovers the group during remote access policy discovery, the group appears in the list under the name <*device_display_name*>DfltGrpPolicy. When you deploy the configuration to the device, the display name prefix is removed so that DfltGrpPolicy is updated correctly. For more information, see Discovering Remote Access VPN Policies, page 30-12.

- To edit an object, select it and click the **Edit Row** button to open the ASA Group Policies Dialog Box, page 34-1.
- To delete an object from the policy, select it and click the **Delete Row** button. The associated policy objects are not deleted, they are only removed from this policy.

  Note You cannot delete the default group policy.

Navigation Path

- (Device view) Select an ASA device, then select **Remote Access VPN > Group Policies** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Group Policies (ASA)** from the Policy selector. Select an existing policy or create a new one.

Understanding Group Policies (ASA)

When you configure a remote access IPSec or SSL VPN connection, you must create user groups to which remote clients will belong. A user group policy is a set of user-oriented attribute/value pairs for remote access VPN connections that are stored either internally (locally) on the device or externally on an AAA server. The connection profile uses a user group policy that sets terms for user connections after the connection is established. Group policies let you apply whole sets of attributes to a user or a group of users, rather than having to specify each attribute individually for each user.

Tip Dynamic Access policies take precedence over Group policies. If a setting is not specified in a Dynamic Access policy, an ASA device checks for Group policies that specify the setting.

An ASA user group comprises the following attributes:

- Group policy source—Identifies whether the user group's attributes and values are stored internally (locally) on the security appliance or externally on an AAA server. If the user group is an external type, no other settings need to be configured for it. For more information, see ASA Group Policies Dialog Box, page 34-1.
- Client Configuration settings, which specify the Cisco client parameters for the user group in an Easy VPN or remote access VPN. For more information, see ASA Group Policies Client Configuration Settings, page 34-6.
- Client Firewall Attributes, which configure the firewall settings for VPN clients in an Easy VPN or remote access VPN. For more information, see ASA Group Policies Client Firewall Attributes, page 34-7.
- Hardware Client Attributes, which configure the VPN 3002 Hardware Client settings in an Easy VPN or remote access VPN. For more information, see ASA Group Policies Hardware Client Attributes, page 34-9.
- IPsec settings, which specify tunneling protocols, filters, connection settings, and servers for the user group in an Easy VPN or remote access VPN. For more information, see ASA Group Policies IPSec Settings, page 34-10.
- Clientless settings, which configure the Clientless mode of access to the corporate network in an SSL VPN, for the ASA user group. For more information, see ASA Group Policies SSL VPN Clientless Settings, page 34-12.
- Full Client settings, which configure the Full Client mode of access to the corporate network in an SSL VPN, for the ASA user group. For more information, see ASA Group Policies SSL VPN Full Client Settings, page 34-19.
- General settings that are required for Clientless/Port Forwarding in an SSL VPN. For more information, see ASA Group Policies SSL VPN Settings, page 34-25.
- DNS/WINS settings that define the DNS and WINS servers and the domain name that should be pushed to remote clients associated with the ASA user group. For more information, see ASA Group Policies DNS/WINS Settings, page 34-30.

- Split tunneling that lets a remote client conditionally direct packets over an IPsec or SSL VPN tunnel in encrypted form or to a network interface in clear text form. For more information, see ASA Group Policies Split Tunneling Settings, page 34-31.
- Remote access or SSL VPN session connection settings for the ASA user group. For more information, see ASA Group Policies Connection Settings, page 34-33.

Related Topics

- Creating Group Policies (ASA, PIX 7.0+), page 31-28
- Configuring Group Policies for Remote Access VPNs, page 31-26

Creating Group Policies (ASA, PIX 7.0+)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the Group Policies page to create group policies for ASA or PIX 7.0+ devices used in remote access IPSec VPNs, or ASA devices used in remote access SSL VPNs. For information about group policies, see:

- Understanding Group Policies (ASA), page 31-27
- Configuring Group Policies for Remote Access VPNs, page 31-26

Step 1 Do one of the following:

- (Device view) With an ASA or PIX 7.0+ device selected, select **Remote Access VPN > Group Policies** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Group Policies (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Group Policies page opens. The table lists the existing group policies, whether they are defined internally on the device or externally on a AAA server, and the protocol for the group: IKEv1 (IPsec), IKEv2 (IPsec), or SSL.

Step 2 Click **Add Row (+)** to open a dialog box from which you can select a user group from a list of predefined ASA user group objects, or create new ones if necessary. To create a new group, click the **Create (+)** button in the dialog box.

Step 3 Select the required ASA user group from the list and click **OK**. If the required group already exists, you are finished.

If the required ASA user group does not exist, create it by clicking **Create (+)**. The Add ASA User Group dialog box appears, displaying a list of settings that you can configure for the ASA user group object. For a description of the elements on this dialog box, see ASA Group Policies Dialog Box, page 34-1.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Select whether to store the ASA user group's attributes and values locally on the device, or on an external server.

Note If you selected to store the ASA user group's attributes on an external server, you do not need to configure any Technology settings. After you specify the AAA server group that will be used for authentication and a password to the AAA server, click **OK** and then select the group in the object selector and click **OK** to add it to the policy.

Step 6 If you selected to store the ASA user group's attributes locally on the device, select the type of VPN for which you are creating the ASA user group from the **Technology** list:

- Easy VPN/IPSec IKEv1—For remote access IPsec VPNs that use IKE version 1 negotiations.
- Easy VPN/IPSec IKEv2—(ASA only.) For remote access IPsec VPNs that use IKE version 2 negotiations.
- SSL Clientless—(ASA only.) For SSL VPNs, all access modes (not just clientless).

Step 7 To configure the user group for Easy VPN/IPSec IKEv1 and Easy VPN/IPSec IKEv2, from the Easy VPN/IPSec VPN folder in the Settings pane:

a. Select **Client Configuration** to configure the Cisco client parameters. For a description of these settings, see ASA Group Policies Client Configuration Settings, page 34-6.

b. Select **Client Firewall Attributes** to configure the firewall settings for VPN clients. For a description of these settings, see ASA Group Policies Client Firewall Attributes, page 34-7.

c. Select **Hardware Client Attributes** to configure the VPN 3002 Hardware Client settings. For a description of these settings, see ASA Group Policies Hardware Client Attributes, page 34-9.

d. Select **IPsec** to specify tunneling protocols, filters, connection settings, and servers. For a description of these settings, see ASA Group Policies IPSec Settings, page 34-10.

Step 8 To configure the user group for an SSL VPN, from the SSL VPN folder in the Settings pane:

a. Select **Clientless** to configure the Clientless mode of access to the corporate network in an SSL VPN. For a description of these settings, see ASA Group Policies SSL VPN Clientless Settings, page 34-12.

b. Select **Full Client** to configure the Full Client mode of access to the corporate network in an SSL VPN. For a description of these settings, see ASA Group Policies SSL VPN Full Client Settings, page 34-19.

c. Select **Settings** to configure the general settings that are required for clientless and thin client (port forwarding) access modes in an SSL VPN. For a description of these settings, see ASA Group Policies SSL VPN Settings, page 34-25.

Step 9 Specify the following settings for an ASA user group in an Easy VPN/IPSec IKEv1 or IKEv2 VPN and SSL VPN configuration, in the Settings pane:

a. Select **DNS/WINS** to define the DNS and WINS servers and the domain name that should be pushed to clients associated with the ASA user group. For a description of these settings, see ASA Group Policies DNS/WINS Settings, page 34-30.

b. Select **Split Tunneling** to allow a remote client to conditionally direct encrypted packets through a secure tunnel to the central site and simultaneously allow clear text tunnels to the Internet through a network interface. For a description of these settings, see ASA Group Policies Split Tunneling Settings, page 34-31.

c. Select **Connection Settings** to configure the SSL VPN connection settings for the ASA user group, such as the session and idle timeouts, including the banner text. For a description of these settings, see ASA Group Policies Connection Settings, page 34-33.

Step 10 Click **OK**.

Step 11 Select the ASA user group from the list and click **OK**.

Understanding SSL VPN Server Verification (ASA)

When connecting to a remote SSL-enabled server through clientless SSL VPN, it is important to know that you can trust the remote server, and that it is in fact the server you are trying to connect to. ASA 9.0 introduces support for SSL server certificate verification against a list of trusted certificate authority (CA) certificates for clientless SSL VPN.

When you connect to a remote server via a web browser using the HTTPS protocol, the server will provide a digital certificate signed by a CA to identify itself. Web browsers ship with a collection of CA certificates which are used to verify the validity of the server certificate. This is a form of public key infrastructure (PKI).

Just as browsers provide certificate management facilities, so does the ASA in the form of trusted certificate pool management facility: trustpools. This can be thought of as a special case of trustpoint representing multiple known CA certificates. The ASA includes a default bundle of certificates, similar to that provided with web browsers, but it is inactive until activated by the administrator.

Note If you are already familiar with trustpools from Cisco IOS then you should be aware that the ASA version is similar, but not identical.

For more information on managing trusted certificates, see the following topics:

- Configuring SSL VPN Server Verification (ASA), page 31-73
- Configuring Trusted Pool Settings (ASA), page 31-30
- Using the Trustpool Manager, page 31-32

Configuring Trusted Pool Settings (ASA)

Use the Trusted Pool Settings page to configure options for certificate revocation. You can also launch the Trustpool Manager.

Navigation Path

(Device View only) Select an ASA device; then select **Remote Access VPN > Trusted Pool** from the Policy selector.

Related Topics

- Configuring SSL VPN Server Verification (ASA), page 31-73
- Using the Trustpool Manager, page 31-324

Field Reference

Table 31-12 Trusted Pool Page

Element	Description
Revocation Check	Whether to check certificates for revocation. Select the appropriate option: • **Check Certificates** If you select this option, also specify the method or methods to use for revocation by selecting the appropriate method (CRL or OCSP) and moving it to the box on the right by clicking **>>**. **Note** You can choose either or both methods. If choosing both methods, add the methods in the order in which you want them used. • **Do not check Certificates**
Certificate Map Settings	Optionally, specify override options for a certificate map by selecting the map from the following lists. Each list will include all certificate maps that are configured on the device. • **Allow Expired Certificates**—Select the certificate map for which you want to allow expired certificates. • **Skip Revocation Check**—Select the certificate map for which you want to skip revocation check.
CRL Options	Specifies options for managing the Certificate Revocation List: • **Cache Refresh Time**—The number of minutes (1-1440) before the ASA considers a CRL too old to be reliable. The default value is 60 minutes. • **Enforce next CRL update**—Whether the ASA should enforce the next CRL update.
Certificate Expiration Alerts	Beginning with version 4.9, Security Manager enables checking all CA and ID certificates in the trust points for expiration once every 24 hours. If a certificate is nearing expiration, a syslog will be issued as an alert. You can configure the reminder and recurrence intervals. This feature is supported only in devices running ASA software version 9.4(1) or higher. **Begin**—Enter the number of days before the expiration in which the first alert will be sent. The range is 1 to 90 days. By default, reminders will start at 60 days prior to expiration. **Repeat**—Enter the frequency, in number of days, at which the alert will be repeated if the certificate is not renewed. The range is 1 to 14 days. By default reminders will be sent every 7 days.

Table 31-12 Trusted Pool Page (continued)

Element	Description
Automatic Import	Beginning with version 4.10, Security Manager enables automatic import of Trustpool certificate bundle. When you enable automatic import, you can configure the URL that the ASA will use to download and import the Trustpool certificate bundle.This feature is supported only in devices running ASA software version 9.5(2) or higher. Beginning with version 4.13, Security Manager provides source interface option that ASA can use to identify the destination URL. This feature is not supported for ASA versions lower to 9.7.1. **Interface**—Click the **Select** button and choose the interface. If the interface configured is management-only, the destination URL is routed through management VRF. For non-management interface, the URL is routed through data VRF. If interface is not specified, both the routing table of the management and data VRF are polled to identify the route to reach the URL. **Import from a URL**—Enter the URL from which the ASA has to download the Trustpool certificate bundle. **Download Time**—Enter the time at which the ASA will download the certificate bundle. The import takes place daily at the time specified here. **Note** The default value of the URL is http://www.cisco.com/security/pki/trs/ios_core.p7b and the default value of the download time is 22:00:00.
Launch Trustpool Manager	Launches the Trustpool Manager, which is used to manage Trustpool certificates. You can use the Trustpool Manager to perform the following: For more information, see Using the Trustpool Manager, page 31-32.

Using the Trustpool Manager

Use the Trustpool Manager to manage the certificates that are included in the trustpool. The Trustpool Manager provides the following functions:

- Updating the trustpool
- Importing a certificate bundle
- Exporting a certificate bundle
- Removing certificates from the trustpool

Navigation Path

(Device View only) Select an ASA device, select **Remote Access VPN > Trusted Pool** from the Policy selector, and then click **Launch Trustpool Manager**.

Updating the Trustpool

The trustpool should be updated if either of the following conditions exists:

- Any certificate in the trustpool is due to expire or has been re-issued.

- The published CA certificate bundle contains additional certificates that are required by a specific application.

To update the certificates in the trustpool, click **Refresh Certificates**.

Importing a Certificate Bundle

You can import individual certificates or bundles of certificates from a variety of locations in one of the following formats:

- x509 certificates in DER format wrapped in a pkcs7 structure
- a file of concatenated x509 certificates in PEM format (complete with PEM header)

To import a certificate or bundle:

1. Click **Import Bundle**.
2. Select the location of the bundle:
   - **Import from Cisco published signed root file distribution**—Select this option to import from the published distribution site.
   - **Import from a URL**—If the bundle is hosted on a server, select this option, select the protocol from the list, and enter the URL in the box.
   - **Bundle file on device**—If the bundle is stored on the ASA flash file system, select this option and then enter the path to the bundle.
   - **Select bundle file**—If the bundle is stored on your machine, click Import from a file, then click Browse Local Files and navigate to the bundle.
   - **Import default bundle**—Select this option to import the default bundle.
3. Specify the following import options:
   - **Clear all certificates before import**—Whether to clear the trustpool before importing the bundle.
   - **Continue to import the bundle if signature validation fails or can't be performed**—Whether to continue import if the signature can not be validated.
4. Click **Import**.

Exporting a Certificate Bundle

When you have correctly configured the Trustpool you should export the pool. This will enable you to restore the Trustpool to this point, for example if you wish to remove a certificate that was added to the trustpool after the export. You can export the pool to the Security Manager server file system or your local file system.

To export the certificate bundle:

1. Click **Export Bundle**.
2. Click **Browse.**
3. Select the tab that corresponds to the file system you want to export to (local machine or Security Manager server).
4. Navigate to the folder where you want to save the trustpool.
5. Enter a unique memorable name for the trustpool in the File name box.
6. Click **Save**.

Removing Certificates from the Trustpool

You can remove certificates from the trustpool using the following methods:

- To remove an individual certificate, select the certificate and click **Delete**.
- To remove all certificates that are not part of the default bundle, click **Clear Trustpool**.

Note Before clearing the trustpool you should export the current trustpool so that you can restore your current settings if needed.

Related Topics

- Configuring SSL VPN Server Verification (ASA), page 31-73
- Configuring Trusted Pool Settings (ASA), page 31-30

Add/Edit Scripts Dialog Box

Use the Add/Edit Scripts dialog box to define a script to use in mapping the username from the certificate.

Navigation Path

- (Device view) Select an ASA device, then select **Remote Access VPN > Username from Cert Scripts** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Username from Cert Scripts (ASA)** from the Policy selector. Select an existing policy or create a new one.

Field Reference

Table 31-13 Add/Edit Scripts Dialog Box

Element	Description
Script Name	Specify the name of the script and use the script in the tunnel group AAA authentication and authorization. The script name may be different for authentication and authorization. You define the script here, and CLI uses the same script to perform this function.
Select Script Parameters	Specify the attributes and content of the script.
Value for Username	Select an attribute from the drop-down list of standard DN attributes to use as the username (Subject DN).
No Filtering	Specify that you want to use the entire specified DN name.
Filter by Substring	Specify the Starting Index (the position in the string of the first character to match) and Ending Index (number of characters to search). If you choose this option, the starting index cannot be blank. If you leave the ending index blank, it defaults to -1, indicating that the entire string is searched for a match.
Filter by Regular Expression	Enter a regular expression to apply to the search in the Regular Expression field. Standard regular expression operators apply.

Table 31-13 Add/Edit Scripts Dialog Box (continued)

Element	Description
Use Custom Script in LUA format	Specify a custom script written in the LUA programming language to parse the search fields. Selecting this option makes available a field in which you can enter your custom LUA script. The following are examples of custom scripts in LUA format: • "return findpattern(cert.subject.cn,"%a+")" • local a,b,c; a,b,c = string.find(cert.subject.fulldn, ',cn=(.+),cn=Users'); return c; **Note** LUA is case-sensitive. The following table provides the attribute names and their descriptions that you can use in an LUA script.

Table 31-14 Attributes in an LUA script

Attribute	Description
cert.subject.c	Country
cert.subject.cn	Common Name
cert.subject.dnq	DN qualifier
cert.subject.ea	E-mail Address
cert.subject.genq	Generational qualified
cert.subject.gn	Given Name
cert.subject.i	Initials
cert.subject.l	Locality
cert.subject.n	Name
cert.subject.o	Organization
cert.subject.ou	Organization Unit
cert.subject.ser	Subject Serial Number
cert.subject.sn	Surname
cert.subject.sp	State/Province
cert.subject.t	Title
cert.subject.uid	User ID
cert.issuer.c	Country
cert.issuer.cn	Common Name
cert.issuer.dnq	DN qualifier
cert.issuer.ea	E-mail Address
cert.issuer.genq	Generational qualified
cert.issuer.gn	Given Name

Table 31-14 Attributes in an LUA script (continued)

Attribute	Description
cert.issuer.i	Initials
cert.issuer.l	Locality
cert.issuer.n	Name
cert.issuer.o	Organization
cert.issuer.ou	Organization Unit
cert.issuer.ser	Issuer Serial Number
cert.issuer.sn	Surname
cert.issuer.sp	State/Province
cert.issuer.t	Title
cert.issuer.uid	User ID
cert.serialnumber	Certificate Serial Number
cert.subjectaltname.upn	User Principal Name

Working with IPSec VPN Policies

Certain policies need to be configured for IPSec VPNs. The topics listed below explain these remote access IPsec VPN policies, with the exception of the IKE Proposal policy, which is explained in Configuring an IKE Proposal, page 26-9.

This section contains the following topics:

- Configuring Certificate to Connection Profile Map Policies (ASA), page 31-36
- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40

Configuring Certificate to Connection Profile Map Policies (ASA)

Certificate to connection profile map policies are used for enhanced certificate authentication on ASA devices in remote access IKEv1 IPSec VPNs. They are not used in remote access IKEv2 IPSec or SSL VPNs.

Certificate to connection profile map policies let you define rules to match a user's certificate to a permission group based on specified fields. To establish authentication, you can use any field of the certificate, or you can have all certificate users share a permission group. You can match the group from the DN rules, the Organization Unit (OU) field, the IKE identity, or the peer IP address. You can use any or all of these methods.

To match user permission groups based on DN fields of the certificate, you define rules that specify the fields to match for a group and then enable each rule for that selected group. A connection profile must already exist in the configuration before you can create a rule for it.

This procedure describes how to configure a Certificate to Connection Profile Map policy for a remote client trying to connect to an ASA server device.

Step 1 Do one of the following:

- (Device View) Select an ASA device; then select **Remote Access VPN > IPSec VPN > Certificate to Connection Profile Maps > Policies** from the Policy selector.
- (Policy View) Select **Remote Access VPN > IPSec VPN > Certificate to Connection Profile Maps > Policies** from the Policy Type selector. Select an existing policy or create a new one.

The Certificate to Connection Profile Map Policies page opens.

Step 2 Select any, or all, of the following options to establish authentication and to determine to which connection profile (tunnel group) to map the client:

- **Use Configured Rules to Match a Certificate to a Group**—To use the rules defined in the Certificate to Connection Profile Maps > Rules policy. For information on configuring the rules, see Configuring Certificate to Connection Profile Map Rules (ASA), page 31-37.
- **Use Certificate Organization Unit (OU) Field to Determine the Group**—To use the organizational unit (OU) field of the client certificate.
- **Use IKE Identify to Determine the Group**—To use the IKE identity.
- **Use Peer IP address to Determine the Group**—To use the peer's IP address.
- **Use Group URL if Group URL and Certificate Map match different Connection profiles** is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Configuring Certificate to Connection Profile Map Rules (ASA)

If you configure certificate to connection profile maps, and select the option to **Use Configured Rules to Match a Certificate to a Group** (as explained in Configuring Certificate to Connection Profile Map Policies (ASA), page 31-36), you need to configure the rules required to match a user to a connection profile based on the user certificate.

To match user permission groups based on fields of the certificate, you define rules that specify the fields to match for a group and then enable each rule for that selected group. You must first define a connection profile (tunnel group) before you can create and map a rule to it.

This procedure describes how to configure the Certificate to Connection Profile Map rules and parameters for any remote client trying to connect to an ASA server device.

Tip Certificate to connection profile map policies apply to remote access IKEv1 IPSec VPNs only. They do not apply to IKEv2 or SSL VPNs.

Before You Begin

- Make sure the connection profiles for which you are creating mapping rules has been configured on the device. See Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.
- Make sure that you select **Use Configured Rules to Match a Certificate to a Group** in the Certificate to Connection Profile Maps Policies policy. See Configuring Certificate to Connection Profile Map Policies (ASA), page 31-36.

Step 1 (Device view only) With an ASA device selected, select **Remote Access VPN > IPSec VPN > Certificate to Connection Profile Maps > Rules** from the Policy selector.

The Certificate to Connection Profile Map Rules page is displayed. The policy has two tables:

- **Maps table (upper table)**—The upper table lists all connection profiles for which you are defining certificate to connection map rules. Each row is a profile map, which includes the name of the connection profile that is being mapped, the priority of the map (lower numbers have higher priority), and the map name. You can configure more than one map for the same connection profile.
  - To configure rules for a map, select it and then use the rules table to create, edit, and delete the rules.
  - To add a map, click the **Add Row** button and fill in the Map Rule Dialog Box (Upper Table), page 31-39.
  - To edit map properties (not rules), select it and click the **Edit Row** button.
  - To delete an entire map, select it and click the **Delete Row** button.
- **Rules table (lower table)**—The rules for the map selected in the upper table. You must ensure that the map is actually selected in the upper table: the group title above the rules table should say "**Details for (Connection Profile Name)**."

  When you select a map, the table shows all rules configured for the map, including the field (subject or issuer), certificate component, matching operator, and the value that the rule is looking for. The remote user must match all configured rules in a map for the device to use the mapped connection profile.
  - To add a rule, click the **Add Row** button and fill in the Map Rule Dialog Box (Lower Table), page 31-39.
  - To edit a rule, select it and click the **Edit Row** button.
  - To delete a rule, select it and click the **Delete Row** button.

Step 2 To add a rule to a map:

a. Select the map in the upper table.

If the map does not already exist, create it by clicking the **Add Row (+)** button beneath the upper table and fill in the Map Rule dialog box for creating maps. In the dialog box, you must select the connection profile for the map, assign a relative priority between 1 and 65535 (lower numbers have higher priority), and a unique map name.

b. Ensure that the map is actually selected. Highlighting the map in the table is not sufficient. The heading above the lower table should be "**Details for (Connection Profile Name)**," and unless the map is new, the table should show some rules.

c. To add a new certificate to connection profile matching rule that must be satisfied in order for a remote client to connect to the device using the profile in this map, click the **Add Row (+)** button beneath the lower table. This opens the Map Rule dialog box with different fields.

Note If you get the error message "Missing Settings, A value ID required for Mapping field, Please select a Mapping," it means that you have not successfully selected a map in the upper table. Click on the desired map again.

d. From the **Field** list, select whether the rule should examine the Subject or Issuer field of the client certificate.

e. From the **Component** list, select the component of the client certificate to use for the matching rule.

f. From the **Operator** field, select how the component should be compared to the Value field: Equals (exact match is required), Contains (the entire value must appear), Does Not Equal, Does Not Contain.

g. In the **Value** field, specify the value to match, then click **OK** to save the rule.

h. Add additional rules to the map as desired.

Step 3 In the **Default Connection Profile** field, select the connection profile that should be used for users who do not meet any of the map rules.

Map Rule Dialog Box (Upper Table)

Use the Map Rule dialog box, when opened for the maps table in the upper part of the Certificate to Connection Profile Maps > Rules policy, to configure maps for which you can then configure rules in the lower table of the Rules policy. For a detailed explanation of configuring these maps and their associated rules, see Configuring Certificate to Connection Profile Map Rules (ASA), page 31-37.

Navigation Path

(Device View only) Select an ASA device; then select **Remote Access VPN > Certificate to Connection Profile Maps > Rules** from the Policy selector. Click the **Add Row** button beneath the upper table, or select a map in the upper table and click **Edit Row**.

Field Reference

Table 31-15 Map Rule Dialog Box (Upper Table)

Element	Description
Map Name	The name of the connection profile map.
Priority	The priority number of the matching rule, between 1 and 65535. A lower number has a higher priority. For example, a matching rule with a priority number of 2, has a higher priority than a matching rule with a priority number of 5. If you create multiple maps, they are processed in priority order, and the first matching rule determines to which profile the user is mapped.
Connection Profile	Select the connection profile for IPSec and for SSL for which you are creating matching rules. You must select either or both connection profiles. Clients attempting to connect to the connection profiles must satisfy the associated matching rule conditions to connect to the device. Connection Profile for IPSec is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Map Rule Dialog Box (Lower Table)

Use the Map Rule dialog box, when opened for the rules table in the lower part of the Certificate to Connection Profile Maps > Rules policy, to configure rules for the map selected in the maps table (upper table of the Rules policy). For a detailed explanation of configuring these rules, see Configuring Certificate to Connection Profile Map Rules (ASA), page 31-37.

Navigation Path

(Device View only) Select an ASA device; then select **Remote Access VPN > IPSec VPN > Certificate to Connection Profile Maps > Rules** from the Policy selector. Click the **Add Row** button beneath the lower table, or select a rule in the lower table and click **Edit Row**.

Field Reference

Table 31-16 Map Rule Dialog Box (Lower Table)

Element	Description
Field	Select the field for the matching rule according to the **Subject** or the **Issuer** of the client certificate.
Component	Select the component of the client certificate to use for the matching rule.
Operator	Select the operator for the matching rule as follows: • Equals—The certificate component must match the entered value. If they do not match exactly, the connection is denied. • Contains—The certificate component must contain the entered value. If the component does not contain the value, the connection is denied. • Does Not Equal—The certificate component *cannot* equal the entered value. For example, for a selected certificate component of Country, and an entered value of US, if the client county value equals US, then the connection is denied. • Does Not Contain—The certificate component *cannot* contain the entered value. For example, for a selected certificate component of Country, and an entered value of US, if the client county value contains US, the connection is denied.
Value	The value of the matching rule. The value entered is associated with the selected component and operator.
Default Connection Profile	This option is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

This procedure describes how to create or edit an IPsec proposal for your remote access VPN server when the server is an ASA or PIX 7.0+ device.

Note Beginning with Cisco Security Manager version 4.17, you can configure and deploy IPSec Proposal policy on ASA multi-context devices running the software version 9.9(2) or later.

If you are configuring an IPsec proposal for IOS or PIX 6.3 devices, including Catalyst 6500/7600 devices, see Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3.

An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA.

When configuring an IPsec proposal, you must define the external interface through which the remote access clients connect to the server, the IKE version to use during IKE negotiation, and the encryption and authentication algorithms that protect the data in the VPN tunnel. You can also enable reverse route injection and NAT traversal.

For more information on IPsec tunnel concepts, see Understanding IPsec Proposals, page 26-18.

Related Topics

- Table Columns and Column Heading Features, page 1-49

Step 1 Do one of the following:

- (Device view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (ASA/PIX 7.x)** from the Policy selector.
- (Policy view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (ASA/PIX 7.x)** from the Policy Type selector. Select an existing policy or create a new one.

The IPsec Proposal page opens and lists the configured proposals, including the VPN endpoint, IPsec transform set, and whether reverse route injection is configured for the proposal.

Step 2 Do any of the following:

- To add a new IPsec proposal, click the **Add Row (+)** button and fill in the IPsec Proposal Editor dialog box. For detailed information on the available options, see IPsec Proposal Editor (ASA, PIX 7.0+ Devices), page 31-41.
- To edit an existing proposal, select it and click the **Edit Row (pencil)** button.
- To delete a proposal, select it and click the **Delete Row (trash can)** button.

IPsec Proposal Editor (ASA, PIX 7.0+ Devices)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

Use the IPsec Proposal Editor to create or edit an IPsec proposal for an ASA or PIX 7.0+ device.

The elements in this dialog box differ according to the selected device. The table below describes the elements on the General tab in the IPsec Proposal Editor dialog box when an ASA or PIX 7.0+ device is selected.

Note For a description of the elements in the dialog box when a PIX 7.0+ or ASA device is selected, see IPsec Proposal Editor (IOS, PIX 6.3 Devices), page 33-4.

Navigation Path

- (Device view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (ASA/PIX 7.x)** from the Policy selector. Click the Add Row (+) or Edit Row (pencil) buttons.

- (Policy view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (ASA/PIX 7.x)** from the Policy Type selector. Select an existing policy or create a new one. Click the Add Row (+) or Edit Row (pencil) buttons.

Related Topics

- Configuring an IPsec Proposal on a Remote Access VPN Server (ASA, PIX 7.0+ Devices), page 31-40
- Understanding IPsec Proposals, page 26-18
- Creating Interface Role Objects, page 6-74
- Creating AAA Server Group Objects, page 6-48

Field Reference

Table 31-17 IPsec Proposal Editor, ASA and PIX 7.0+ Devices)

Element	Description
External Interface	The external interface through which remote access clients will connect to the server. Enter the name of the interface or interface role object, or click **Select** to select it or to create a new object.
Enable IKEv1 Enable IKEv2	The IKE versions to use during IKE negotiations. IKEv2 is supported on ASA Software release 8.4(1)+ only with Anyconnect 3.0+ clients. Select either or both options as appropriate.
Enable Client Services Client Services Port Number	Available only if you enable IKEv2. Whether to enable the Client Services Server on the ASA for this connection. The Client Services Server provides HTTPS (SSL) access to allow the AnyConnect Downloader to receive software upgrades, profiles, localization and customization files, CSD, SCEP, and other file downloads required by the AnyConnect client. If you select this option, specify the client services port number, which is 443 by default. If you do not enable the Client Services Server, users will not be able to download any of these files that the AnyConnect client might need. **Tip** You can use the same port that you use for SSL VPN running on the same device. Even if you have an SSL VPN configured, you must select this option to enable file downloads over SSL for IKEv2 IPsec clients.
IKEv1 Transform Sets IKEv2 Transform Sets	The transform sets to use for your tunnel policy. Transform sets specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel. The transform sets are different for each IKE version; select objects for each supported version. You can select up to 11 transform sets for each. For more information, see Understanding Transform Sets, page 26-20. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security will be used. Click **Select** to select the IPsec transform set policy objects to use in the topology. If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27.

Table 31-17 IPsec Proposal Editor, ASA and PIX 7.0+ Devices) (continued)

Element	Description
Reverse Route Injection	Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint. For more information, see Understanding Reverse Route Injection, page 26-21. Select one of the following options to configure RRI on the crypto map: • **None**—Disables the configuration of RRI on the crypto map. • **Standard**—Creates routes based on the destination information defined in the crypto map access control list (ACL). This is the default option.
Enable Network Address Translation Traversal	Whether to allow Network Address Translation traversal (NAT-T). Use NAT traversal when there is a device between a VPN-connected hub and spoke that performs Network Address Translation (NAT) on the IPsec traffic. For information about NAT traversal, see Understanding NAT in VPNs, page 26-41
ESPv3 Settings (ASA 9.0.1+ only) Specify whether incoming ICMP error messages are validated for cryptography and dynamic cryptography maps, set the per-security association policy, or enable traffic flow packets:	
Validate Incoming ICMP error messages	Whether to validate those ICMP error messages received through an IPsec tunnel and destined for an interior host on the private network.
Enable Do Not Fragment (DF) Policy	Define how the IPsec subsystem handles large packets that have the do-not-fragment (DF) bit set in the IP header. Choose one of the following: • Set—Sets and uses the DF bit. • Copy—Maintains the DF bit. • Clear—Ignores the DF bit.
Enable Traffic Flow Confidentiality (TFC) Packets	Enable dummy TFC packets that mask the traffic profile which traverses the tunnel. **Note** You must have an IKEv2 IPsec proposal set on the Tunnel Policy (Crypto Map) Basic tab before enabling TFC. Traffic Flow Confidentiality is not available when IKEv1 is enabled. Use the Burst, Payload Size, and Timeout parameters to generate random length packets at random intervals across the specified SA.

Working with SSL and IKEv2 IPSec VPN Policies

Certain policies need to be configured for SSL VPNs. These policies are also used with remote access IKEv2 IPSec VPNs. The topics listed below explain these remote access VPN policies.

This section contains the following topics:

- Understanding SSL VPN Access Policies (ASA), page 31-44
- Configuring Other SSL VPN Settings (ASA), page 31-51

- Configuring SSL VPN Shared Licenses (ASA 8.2+), page 31-74

Understanding SSL VPN Access Policies (ASA)

An Access policy specifies the security appliance interfaces on which a remote access SSL or IKEv2 IPsec VPN connection profile can be enabled, the port to be used for the connection profile, Datagram Transport Layer Security (DTLS) settings, the SSL VPN session timeout and maximum number of sessions. You can also specify whether to use the AnyConnect VPN Client or AnyConnect Essentials Client.

For more information about the Anyconnect VPN Client, see Understanding SSL VPN AnyConnect Client Settings, page 31-62. The remainder of this topic explains DTLS and AnyConnect Essentials in more detail.

Datagram Transport Layer Security (DTLS)

Enabling Datagram Transport Layer Security (DTLS) allows the AnyConnect client establishing an SSL VPN connection to use two simultaneous tunnels—an SSL tunnel and a DTLS tunnel. Using DTLS avoids latency and bandwidth problems associated with SSL connections and improves the performance of real-time applications that are sensitive to packet delays. By default, DTLS is enabled when SSL VPN access is enabled on an interface. If you disable DTLS, SSL VPN connections connect with an SSL VPN tunnel only.

Note In order for DTLS to fall back to a TLS connection, you must specify a fallback trustpoint. If you do not specify a fallback trustpoint and the DTLS connection experiences a problem, the connection terminates instead of falling back to the specified trustpoint.

AnyConnect Essentials VPN Client

AnyConnect Essentials is a separately licensed VPN client for SSL or IKEv2 IPsec, entirely configured on the adaptive security appliance, that provides the full AnyConnect capability, with the following exceptions:

- No CSD (including HostScan/Vault/Cache Cleaner)
- No clientless SSL VPN
- Optional Windows Mobile Support

The AnyConnect Essentials client provides remote end users running Microsoft Windows Vista, Windows Mobile, Windows XP or Windows 2000, Linux, or Macintosh OS X, with the benefits of a Cisco VPN client. If this feature is disabled, the full AnyConnect VPN client is used. This feature is disabled by default.

Note This license cannot be used at the same time as the shared license for SSL VPN.

This section contains the following topics:

- SSL VPN Access Policy Page, page 31-45
- Configuring an Access Policy, page 31-49

SSL VPN Access Policy Page

Use the SSL VPN Access Policy page to configure access parameters for your remote access SSL or IKEv2 IPsec VPN. For more information about configuring an Access policy, see Configuring an Access Policy, page 31-49.

Tip Any trustpoints that you specify in this policy must also be selected in the **Public Key Infrastructure** policy. For more information, see Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.

Navigation Path

- (Device View) Select **Remote Access VPN > SSL VPN > Access** from the Policy selector.
- (Policy View) Select **Remote Access VPN > SSL VPN > Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Related Topics

- Understanding SSL VPN Access Policies (ASA), page 31-44
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 31-18 SSL VPN Access Policy Page

Element	Description
Access Interface Table	The Access Interface table lists the interfaces that are configured for remote access SSL or IKEv2 IPSec VPN connections. The table displays the access settings for each interface: whether the interface is enabled to allow VPN access, whether DTLS is enabled, whether client certificates are required, and the trustpoints used by the interface. • To configure access on an interface, click the **Add row (+)** button (see Access Interface Configuration Dialog Box, page 31-48). • To edit access settings for an interface, select the interface and click the **Edit Row (pencil)** button (see Access Interface Configuration Dialog Box, page 31-48). • To delete access settings for an interface, select the interface and click the **Delete Row (trash can)** button.

Table 31-18 SSL VPN Access Policy Page (continued)

Element	Description
Server Name Indication Table	The Server Name Indication table lists the Server Name Indication mappings that have been defined. • To define a Server Name Indication mapping, click the **Add row (+)** button (see Server Name Indication Dialog Box, page 31-49). • To edit an existing mapping, select the mapping and click the **Edit Row (pencil)** button (see Server Name Indication Dialog Box, page 31-49). • To delete a Server Name Indication mapping, select the mapping and click the **Delete Row (trash can)** button. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Port Number	The port to use for VPN sessions. The default port is 443, for HTTPS traffic. If HTTP port redirection is enabled, the default HTTP port number is 80. To specify a non-default port, the range is 1024 through 65535. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode. Enter the port number or the name of a port list object, or click **Select** to select a port list object or to create a new object. **Note** If you change the port number, all current SSL VPN connections terminate (upon configuration deployment), and current users must reconnect.
DTLS Port Number	The UDP port to use for DTLS connections. The default port is 443. For details about DTLS, see Understanding SSL VPN Access Policies (ASA), page 31-44. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode. Enter the port number or the name of a port list object, or click **Select** to select a port list object or to create a new object.
Fallback Trustpoint	The trustpoint (Certificate Authority, or CA server) to use for interfaces that do not have an assigned trustpoint. Enter the name of a PKI enrollment object, or click **Select** to select the object from a list or to create a new object. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Table 31-18 SSL VPN Access Policy Page (continued)

Element	Description
Default Idle Timeout	The amount of time, in seconds, that an SSL or IKEv2 IPSec VPN session can be idle before the security appliance terminates it. This value applies only if the Idle Timeout value in the group policy for the user is set to zero (0), which means there is no timeout value; otherwise the group policy Idle Timeout value takes precedence over the timeout you configure here. The minimum value you can enter is 60 seconds (1 minute). The default is 30 minutes (1800 seconds). The maximum is 24 hours (86400 seconds). We recommend that you set this attribute to a short time period. This is because a browser set to disable cookies (or one that prompts for cookies and then denies them) can result in a user not connecting but nevertheless appearing in the sessions database. If the Simultaneous Logins attribute for the group policy is set to one, the user cannot log back in because the database indicates that the maximum number of connections already exists. Setting a low idle timeout removes such phantom sessions quickly, and lets a user log in again. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Max Session Limit	The maximum number of SSL or IKEv2 IPSec VPN sessions allowed. Be aware that the different ASA models have different maximum session limits: • ASA 5505—25. • ASA 5510—250. • ASA 5520—750. • ASA 5540—2500. • ASA 5550, 5585-X with SSP-10—5000. • ASA 5580, 5585-X (other models)—10,000. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Certificate Authentication Timeout (ASA 8.4(5) or ASA 9.1(2)+)	The amount of time, in minutes, to wait before timing out certificate authentication. Valid values are from 1 to 120 minutes. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Allow Users to Select Connection Profile in Portal Page	Whether to present a list of configured connection profiles (tunnel groups) from which the user can select the appropriate profile when the user logs in (for example, in the SSL VPN portal page). If you do not select this option, the user cannot select a profile and must use the default profile for the connection. **Tip** You must select this option for remote access IKEv2 IPSec VPNs. It is optional for SSL VPNs. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Table 31-18 SSL VPN Access Policy Page (continued)

Element	Description
Enable AnyConnect Access	Whether to allow the user to use the AnyConnect VPN client to make an SSL or IKEv2 IPSec VPN connection. The option is selected by default. For details about AnyConnect VPN clients, see Understanding SSL VPN AnyConnect Client Settings, page 31-62. **Tip** You must select this option for remote access IKEv2 IPSec VPNs. For SSL VPN, select this option if you want to enable full client access. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Enable AnyConnect Essentials	Whether to enable the AnyConnect Essentials feature, which can be used with both SSL and IKEv2 IPSec VPNs. For details about AnyConnect Essentials VPN clients, see Understanding SSL VPN Access Policies (ASA), page 31-44.

Access Interface Configuration Dialog Box

Use the Access Interface Configuration dialog box to configure an interface on an ASA device for remote access SSL or IKEv2 IPSec VPN connections.

Navigation Path

Open the SSL VPN Access policy (see SSL VPN Access Policy Page, page 31-45), then click **Add Row** below the interface table, or select a row in the table and click **Edit Row**.

Related Topics

- Configuring an Access Policy, page 31-49
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 31-19 Access Interface Configuration Dialog Box

Element	Description
Access Interface	The interface or interface role object on which you want to configure SSL or IKEv2 IPSec VPN access. Enter the name of the interface or interface role, or click **Select** to select one from a list or to create new interface role objects. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Trustpoint Load Balancing Trustpoint	The trustpoint (Certificate Authority, or CA server) to use for authenticating users on the interface. Enter the name of a PKI enrollment object, or click **Select** to select one or to create a new object. If load balancing is configured, you can also select a separate PKI enrollment object for the load balancing trustpoint. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Table 31-19 Access Interface Configuration Dialog Box (continued)

Element	Description
Allow Access	Select this option to enable VPN access via this interface. If the option is not selected, access is configured on the interface, but it is disabled. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
Enable DTLS	When selected, enables Datagram Transport Layer Security (DTLS) on the interface and allows an AnyConnect VPN Client to establish an SSL VPN connection using two simultaneous tunnels—an SSL tunnel and a DTLS tunnel. This is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Server Name Indication Dialog Box

Beginning from version 4.8, Security Manager enables you to configure Server Name Indication mappings used by the enabled VPN interface for authentication. This capability includes the mapping of domain names to trustpoints.

Use the Server Name Indication dialog box to define or modify a domain and trustpoint for each interface.

NOTES:

- You can configure a unique domain name to a trustpoint. However, you can map a trustpoint to multiple domain names. You can configure a maximum of 16 unique trustpoints.
- Server Name Indication mapping of domain names to trustpoints is supported for devices that are running the ASA software version 9.3(2) or higher.

Navigation Path

Open the SSL VPN Access policy (see SSL VPN Access Policy Page, page 31-45), then click **Add Row** below the ServerNameIndication table, or select a row in the table and click **Edit Row**.

Field Reference

Table 31-20 Server Name Indication Dialog Box

Element	Description
Domain Mask	Enter the domain name that the trustpoint will be configured with. This domain will not be associated with any particular interface. A certificate with an associated domain may be used by any interface.
Trustpoint	The trustpoint (Certificate Authority, or CA server) to use for authenticating users on the interface. Enter the name of a PKI enrollment object, or click **Select** to select one or to create a new object.

Configuring an Access Policy

This procedure describes how to configure an Access policy on an ASA device. Access policies are required for remote access SSL and IKEv2 IPSec VPN connections. For more information about access policies, see Understanding SSL VPN Access Policies (ASA), page 31-44.

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Access** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Access page opens. For a description of the elements on this page, see SSL VPN Access Policy Page, page 31-45.

Step 2 In the interface table at the top of the policy, configure all of the interfaces on which you will allow remote access SSL or IKEv2 IPSec VPN connections:

- To add an interface, click the **Add Row (+)** button beneath the table and fill in the Add Access Interface Configuration dialog box. You must specify the interface name (or an interface role object that identifies the desired interfaces) and whether to allow access on the interface.

  You can also specify the PKI enrollment object that identifies the Certificate Authority (CA) server trustpoint for the interface (and a load balancing trustpoint if you use load balancing), whether to enable DTLS connections, and whether to require that the client have a valid certificate to complete a connection. For details about the options, see Access Interface Configuration Dialog Box, page 31-48.
- To edit the settings for an interface, select it and click the **Edit Row (pencil)** button.
- To delete an interface, select it and click the **Delete Row** button. Keep in mind that you can edit the interface settings to disable access, so you should delete an interface only if you want to permanently remove it from VPN use.

Step 3 Configure the remaining settings. The settings are described in detail in SSL VPN Access Policy Page, page 31-45. The following are the settings that are of particular interest:

- **Fallback Trustpoint**—The Certificate Authority (CA) server trustpoint to use if an interface does not have a trustpoint configured in the table. Enter the name of a PKI enrollment object, or click **Select** to select one or to create a new object.
- **Allow Users to Select Connection Profile in Portal Page**—If you have multiple tunnel groups, selecting this option allows the user to select the correct tunnel group during login. You must select this option for IKEv2 IPSec VPNs.
- **Enable AnyConnect Access**—The AnyConnect VPN client is a full client; you must enable AnyConnect access if you want to allow full client access to the VPN. You must select this option for IKEv2 IPSec VPNs.

  For more information about AnyConnect, including AnyConnect Essentials, see Understanding SSL VPN AnyConnect Client Settings, page 31-62.
- **Enable AnyConnect Essentials**—Select this option if you are using AnyConnect Essentials clients, which you can use with remote access SSL or IKEv2 IPSec VPNs.

Step 4 Any trustpoints that you specify in this policy must also be selected in the **Public Key Infrastructure** policy. For more information, see Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.

Configuring Other SSL VPN Settings (ASA)

The SSL VPN Other Settings policy for ASA devices defines settings that include caching, content rewriting, character encoding, proxy and proxy bypass definitions, browser plug-ins, AnyConnect client images and profiles, Kerberos Constrained Delegation, and some other advanced settings.

To configure the Other Settings policy, do one of the following:

- (Device View) Select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy View) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

You can then configure the settings on the following tabs:

- Performance tab—To configure caching to improve SSL VPN performance. See Configuring SSL VPN Performance Settings (ASA), page 31-52. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Content Rewrite tab—To create rules that permit users to browse certain sites and applications without going through the security appliance itself. See Configuring SSL VPN Content Rewrite Rules (ASA), page 31-53. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Encoding tab—To configure non-default encoding for web pages delivered from CIFS servers. Encoding is normally determined by the remote user's browser. See Configuring SSL VPN Encoding Rules (ASA), page 31-55. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Proxy tab—To define HTTP or HTTPS proxy servers, if your network requires them, and proxy bypass rules. See Configuring SSL VPN Proxies and Proxy Bypass (ASA), page 31-57. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Plug In tab—To define browser plug-ins, which are separate programs that a web browser invokes to perform a dedicated function. See Configuring SSL VPN Browser Plug-ins (ASA), page 31-60. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Client Settings tab—To configure AnyConnect client images and profiles for downloading to clients. See the following topics:
  - Understanding SSL VPN AnyConnect Client Settings, page 31-62
  - Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64

  This is partially supported for ASA 9.5(2) Remote Access VPN in Multi-context mode. Only the AnyConnect Client Image is supported in ASA 9.5(2) Multiple Context Mode. Beginning with version 4.12, Security Manager provides support for multi-context ASA 9.6(2) and later devices for Admin and User contexts. The CLIs supported are:
  - Anyconnect image
  - Anyconnect profile

  During discovery, the AnyConnect Image is not discovered for ASA 9.5(2) remote access VPN Multiple Context mode. After discovery if you want to remove the AnyConnect Image configuration you must use FlexConfig.
- Microsoft KCD Server—To configure Kerberos Constrained Delegation (KCD) for use with clientless SSL VPN connections. See the following topics:
  - Understanding Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-66
  - Configuring Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-69

This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

- AnyConnect Custom Attributes tab—To configure AnyConnect custom attributes. See Configuring AnyConnect Custom Attributes (ASA), page 31-70. AnyConnect Custom Attributes tab is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.
- Advanced tab—To configure the memory, on-screen keyboard, and internal password features. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Note Beginning with 4.15, Cisco Security Manager supports HTTP Strict Transport Security (HSTS). HTTS is a web security policy mechanism which helps to protect websites against protocol downgrade attacks and hijack of cookies.

You can enable or disable HSTS and also provide the timeout values in the Advanced tab. See Configuring SSL VPN Advanced Settings (ASA), page 31-72.

- SSL Server Verification tab—To enable HTTPS server verification for clientless SSL VPN users. See Configuring SSL VPN Server Verification (ASA), page 31-73. This is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Tip You must also configure a connection profile policy on the device. See Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.

Configuring SSL VPN Performance Settings (ASA)

Caching enhances SSL VPN performance. It stores frequently reused objects in the system cache, which reduces the need to perform repeated rewriting and compressing of content. It reduces traffic between SSL VPN and both the remote servers and end-user browsers, with the result that many applications run much more efficiently.

This procedure describes how to enable caching on your ASA security appliance.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector. Click the **Performance** tab if it is not already selected.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one. Click the **Performance** tab if it is not already selected.

Step 2 Select **Enable** to enable caching on the security appliance.

If you deselect this option, the cache settings configured on the security appliance do not take effect.

Step 3 Configure the following options:

- **Minimum Object Size**—The minimum size of an HTTP object that can be stored in the cache on the security appliance, in kilobytes. The range is 0-10,000 KB. The default is 0 KB.
- **Maximum Object Size**—The maximum size of an HTTP object that can be stored in the cache on the security appliance, in kilobytes. The range is 0-10,000 KB. The default is 1000 KB. The maximum size must be larger than the minimum size.

- **Last Modified Factor**—An integer to set a revalidation policy for caching objects that have only the last-modified timestamp, and no other server-set expiration values. The range is 1-100. The default is 20.

  The Expires response from the origin web server to the security appliance request, which indicates the time that the response expires, also affects caching. This response header indicates the time that the response becomes stale and should not be sent to the client without an up-to-date check (using a conditional GET operation).

  The security appliance can also calculate an expiration time for each web object before it is written to disk. The algorithm to calculate an object's cache expiration date is as follows:

  Expiration date = (Today's date - Object's last modified date) * Freshness factor

  After the expiration date has passed, the object is considered stale and subsequent requests causes a fresh retrieval of the content by the security appliance. Setting the last modified factor to zero is equivalent to forcing an immediate revalidation, while setting it to 100 results in the longest allowable time until revalidation.
- **Expiration Time**—The amount of time (in minutes) that the security appliance caches objects without revalidating them. The range is 0-900 minutes. The default is one minute.

  Revalidation consists of rejecting the objects from the origin server before serving the requested content to the client browser when the age of the cached object has exceeded its freshness lifetime. The age of a cached object is the time that the object has been stored in the security appliance's cache without the security appliance explicitly contacting the origin server to check if the object is still fresh.
- **Cache Static Content**—Whether to cache static content on the security appliance. Each web page can include static and dynamic objects. The security appliance caches individual static objects, such as image files (*.gif, *.jpeg), java applets (.js), and cascading style sheets (*.css).

Configuring SSL VPN Content Rewrite Rules (ASA)

SSL VPN processes application traffic through a content transformation/rewriting engine that includes advanced elements (such as, JavaScript, VBScript, Java, and multi-byte characters) to proxy HTTP traffic depending on whether the user is using an application within or independently of an SSL VPN device.

If you do not want some applications and web resources, such as public web sites, to go through the security appliance, you can create rewrite rules that permit users to browse certain sites and applications without going through the security appliance itself. This is similar to split tunneling in an IPsec VPN connection.

In the Content Rewrite tab of the SSL VPN Other Settings page, you can configure multiple content rewrite rules. The Content Rewrite tab lists all applications for which content rewrite is enabled or disabled.

Tip The security appliance searches rewrite rules by order number, starting with the lowest, and applies the first rule that matches.

This procedure shows you how to create or edit content rewrite rules.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Content Rewrite** tab. The Content Rewrite tab displays all applications for which content rewrite is enabled or disabled.

The security appliance searches rewrite rules by order number, starting with the lowest, and applies the first rule that matches. The resource mask defines the application string to which the rule is matched.

If a rule does not have a number, it is evaluated after all of the numbered rules.

Step 3 Do any of the following:

- To add a rule, click the **Add Row** button beneath the table and fill in the Add Content Rewrite dialog box. The options are described in detail in Add/Edit Content Rewrite Dialog Box, page 31-54.
- To edit a rule, select it, click the **Edit Row** button, and make your changes in the Edit Content Rewrite dialog box.
- To delete a rule, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Add/Edit Content Rewrite Dialog Box

Use the Add or Edit Content Rewrite dialog box to configure the rewriting engine that includes advanced elements such as JavaScript, VBScript, Java, and multi-byte characters to proxy HTTP traffic over a SSL VPN connection. For more information about content rewrite rules, see Configuring SSL VPN Content Rewrite Rules (ASA), page 31-53.

Navigation Path

From the Content Rewrite tab of the SSL VPN Other Settings policy for ASA devices, click the **Add Row** button, or select a rule and click the **Edit Row** button. For detailed information on opening the tab, see Configuring SSL VPN Content Rewrite Rules (ASA), page 31-53.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Field Reference

Table 31-21 Add or Edit Content Rewrite Dialog Box

Element	Description
Enable	When selected, enables content rewriting on the security appliance for the rewrite rule. Some applications do not require this processing, such as external public web sites. For these applications, you might choose to turn off content rewriting.

Table 31-21 Add or Edit Content Rewrite Dialog Box (continued)

Element	Description
Rule Number	The number for this rule. This number specifies the position of the rule in the list. Rules without a number are at the end of the list. The range is from 1 to 65534. Rules are processed from the lowest to the highest number, and the first match is applied to the traffic.
Rule Name	An alphanumeric string that describes the content rewrite rule. The maximum length is 128 characters.
Resource Mask	The name of the application or resource to which the rule applies. The maximum length is 300 characters. You can use the following wildcards: • *—Matches everything. You cannot use this wildcard by itself. It must accompany an alphanumeric string. • ?—Matches any single character. • [!x-y]—Matches any character not in the sequence. • [x-y]—Matches any character in the sequence.

Configuring SSL VPN Encoding Rules (ASA)

Use the Encoding tab of the SSL VPN Other Settings page to specify the character set to encode in SSL VPN portal pages to be delivered to remote users. By default, the encoding type set on the remote browser determines the character set for SSL VPN portal pages, so you need to set the character encoding only if it is necessary to ensure proper encoding on the browser.

Character encoding is the pairing of raw data (such as 0's and 1's) with characters to represent the data. The language determines the character encoding method to use. Some languages use the same method, while others do not. Usually, the geographic region determines the default encoding method used by the browser, but the remote user can change this. The browser can also detect the encoding specified on the page, and render the document accordingly.

The encoding attribute lets you specify the value of the character encoding method in the SSL VPN portal page to ensure that the browser renders it properly, regardless of the region in which the user is using the browser, or any changes made to the browser.

The character encoding attribute is a global setting that, by default, all SSL VPN portal pages inherit. However, you can override the file-encoding attribute for Common Internet File System (CIFS) servers that use character encoding that differs from the value of the character-encoding attribute. You can use different file-encoding values for CIFS servers that require different character encodings.

The SSL VPN portal pages downloaded from the CIFS server to the SSL VPN user encode the value of the SSL VPN file-encoding attribute identifying the server, or if one does not, they inherit the value of the character encoding attribute. The remote user's browser maps this value to an entry in its character encoding set to determine the proper character set to use. The SSL VPN portal pages do not specify a value if SSL VPN configuration does not specify a file encoding entry for the CIFS server and the character encoding attribute is not set. The remote browser uses its own default encoding if the SSL VPN portal page does not specify the character encoding, or if it specifies a character encoding value that the browser does not support.

In the Encoding tab of the SSL VPN Global Settings page, you can view the currently configured character sets associated with the CIFS server to be encoded in the portal pages. From this tab, you can create or edit the character sets, as described in the following procedure.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Encoding** tab. The Encoding tab displays the default encoding and a list of CIFS servers for which encoding rules are configured.

Step 3 From the **Global SSL VPN Encoding Type** list, select the attribute that determines the character encoding that all SSL VPN portal pages inherit, except for those from the CIFS servers listed in the table.

Note If you choose **none** or specify a value that the browser on the SSL VPN client does not support, the browser uses its own default encoding. The default global encoding is none.

You can select from the following encoding types:

- big5
- gb2312
- ibm-850
- iso-8859-1
- shift_jis

Note If you are using Japanese Shift_jis Character encoding, click **Do not specify** in the Font Family area of the associated Select Page Font pane to remove the font family.

- unicode
- windows-1252
- none

Step 4 Do any of the following:

- To add a rule, click the **Add Row** button beneath the table and configure the following settings in the Add File Encoding dialog box:
  - **CIFS Server IP, CIFS Server Host**—Select one of these options to specify the CIFS server either by IP address or hostname. If you select IP address, you can either enter the IP address or the name of a network/host object that specifies one or more individual IP addresses.

    If you specify a hostname, the security appliance retains the case you specify, although it ignores the case when matching the name to a server.
  - **Encoding Type**—Select the encoding type. The options are the same as for the global setting described above.

- To edit a rule, select it, click the **Edit Row** button, and make your changes in the Edit File Encoding dialog box.
- To delete a rule, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Configuring SSL VPN Proxies and Proxy Bypass (ASA)

Use the Proxy tab of the SSL VPN Other Settings page to configure the security appliance to terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers. On this tab, you can also configure the security appliance to perform minimal content rewriting and to specify the types of content to rewrite—external links, XML, or neither.

The security appliance can terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers. These servers act as intermediaries between users and the Internet. Requiring all Internet access through a server you control provides another opportunity for filtering to assure secure Internet access and administrative control.

Note The HTTP/HTTPS proxy does not support connections to personal digital assistants.

You can specify a proxy auto-configuration (PAC) file to download from an HTTP proxy server; however, you cannot use proxy authentication when specifying the PAC file.

You can configure the security appliance to use proxy bypass when applications and web resources work better with the content rewriting this feature provides. Proxy bypass is an alternative method of content rewriting that makes minimal changes to the original content. It is useful with custom web applications.

You can configure multiple proxy bypass entries. The order in which you configure them is unimportant. The interface and path mask or interface and port uniquely identify a proxy bypass rule.

If you configure proxy bypass using ports rather than path masks, depending on your network configuration, you might need to change your firewall configuration to allow these ports access to the security appliance. Use path masks to avoid this restriction. Be aware, however, that path masks can change, so you might need to use multiple path mask statements to exhaust the possibilities.

This procedure shows you how to define proxies and proxy bypass rules for your SSL VPN.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Proxy** tab. The Proxy tab displays any currently defined proxies and proxy rules.

Step 3 From the **Proxy Type** field, select the type of external proxy server to use for SSL VPN connections:

- **HTTP/HTTPS Proxy Server**—To specify proxy servers to handle HTTP or HTTPS requests.

- **Proxy Using PAC**—To specify a proxy auto-configuration (PAC) file to download from an HTTP proxy server to the user's browser. Once downloaded, the PAC file uses a JavaScript function to identify a proxy for each URL.

  If you select this option, enter the URL for the PAC file in the **Specify Proxy Auto Config file URL** field. The URL must begin with **http://** or the security appliance will not use the PAC file.

Step 4 If you select HTTP/HTTPS Proxy Server for the proxy type, configure the settings for the HTTP and HTTPS proxy servers. There are separate settings for the HTTP and HTTPS server, allowing you to use different servers, or to specify only one type of proxy. Configure the following options:

- **Enable HTTP Proxy Server, Enable HTTPS Proxy Server**—Select either or both of these options to configure the proxy server.
- **HTTP Proxy Server (IPv4/IPv6), HTTPS Proxy Server (IPv4/IPv6)**—Enter the IP address, or the name of a network/host object that contains the single proxy server's IP address, for each type of proxy server you are configuring. You can click **Select** to select the object from a list or to create a new object.

  The default ports are 80 for HTTP and 443 for HTTPS.

  Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA 9.0(1) or later devices. If the IPv6 address you entered is invalid, Security Manager would throw an error. Security Manager displays a warning message if the object is not available when you select the proxy server from the list.
- **HTTP Proxy Port, HTTPS Proxy Port**—Enter the port on the proxy server to which HTTP or HTTPS requests will be forwarded. You can also enter the name of a port list object that defines the port, or click **Select** to select an object or to create a new one.
- **Exception Address List**—A URL or a comma-delimited list of several URLs to exclude from those that should be sent to the HTTP or HTTPS proxy servers. The string does not have a character limit, but the entire command cannot exceed 512 characters. You can specify literal URLs or use the following wildcards:
  - ***** to match any string, including slashes (/) and periods (.). You must accompany this wildcard with an alphanumeric string.
  - **?** to match any single character, including slashes and periods.
  - **[x-y]** to match any single character in the range of x and y, where x represents one character and y represents another character in the ANSI character set.
  - **[!x-y]** to match any single character that is not in the range.
- **Authentication User Name, Authentication Password, Confirm**—If the proxy server requires user authentication, enter a valid user name and password.

Step 5 If necessary, configure proxy bypass rules in the Proxy Bypass table at the bottom of the tab. Proxy bypass rules specify the ASA interface, port, and target URL configured for proxy bypass. Do any of the following:

- To create a proxy bypass rule, click the **Add Row** button and fill in the Add Proxy Bypass dialog box. For specific information on the attributes of a proxy bypass rule, see Add or Edit Proxy Bypass Dialog Box, page 31-59.
- To edit a proxy bypass rule, select the rule and click the **Edit Row** button.
- To delete a rule, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Tip If you configure proxy bypass rules, you must also configure the SSL VPN Access policy. For more information, see Configuring an Access Policy, page 31-49.

Add or Edit Proxy Bypass Dialog Box

Use the Add or Edit Proxy Bypass dialog box to set proxy bypass rules when the security appliance should perform little or no content rewriting.

Navigation Path

From the Proxy tab of the SSL VPN Other Settings policy for ASA devices, click the **Add Row** button, or select a rule and click the **Edit Row** button. For detailed information on opening the tab, see Configuring SSL VPN Encoding Rules (ASA), page 31-55.

Field Reference

Table 31-22 Add or Edit Proxy Bypass Dialog Box

Element	Description
Interface	The interface on the security appliance that is used for proxy bypass. Enter the name of the interface or the interface role object, or click **Select** to select it from a list or to create a new object.
Bypass On Port	Select this option to use a port number for proxy bypass. Valid port numbers are 20000-21000. Enter the ports or the name of a port list object, or click **Select** to select an object or to create a new one. **Note** If you configure proxy bypass using ports rather than path masks, depending on your network configuration, you might need to change your firewall configuration to allow these ports access to the security appliance. Use path masks to avoid this restriction.
Bypass Matching Specific Pattern	Select this option to use a URL path mask to match for proxy bypass. A path is the text in a URL that follows the domain name. For example, in the URL www.mycompany.com/hrbenefits, *hrbenefits* is the path. You can use the following wildcards: • *****—Matches everything. You cannot use this wildcard by itself. It must accompany an alphanumeric string. • **?**—Matches any single character. • **[x-y]**—Matches any character in the sequence. • **[!x-y]**—Matches any character not in the sequence. The maximum is 128 bytes. **Note** Path masks can change, so you might need to use multiple path mask statements to exhaust the possibilities.

Table 31-22 Add or Edit Proxy Bypass Dialog Box (continued)

Element	Description
URL	Select the **http** or **https** protocol, then enter a URL to which you want to apply proxy bypass. URLs used for proxy bypass allow a maximum of 128 bytes. The port for HTTP is 80 and for HTTPS it is 443, unless you specify another port.
Rewrite XML	Whether to rewrite XML sites and applications to be bypassed by the security appliance.
Rewrite Hostname	Whether to rewrite external links to be bypassed by the security appliance.

Configuring SSL VPN Browser Plug-ins (ASA)

A browser plug-in is a separate program that a web browser invokes to perform a dedicated function, such as connect a client to a server within the browser window. The security appliance lets you import plug-ins for download to remote browsers in clientless SSL VPN sessions.

Cisco redistributes the following open-source, Java-based components to be accessed as plug-ins for web browsers in Clientless SSL VPN sessions. Cisco tests the plug-ins it redistributes, and in some cases, tests the connectivity of plug-ins we cannot redistribute. These files are available in the \files\vms\repository folder in the product installation folder (usually C:\Program Files\CSCOpx) on the Security Manager server. The actual file names include release numbers:

- rdp-plugin.jar—The Remote Desktop Protocol plug-in lets the remote user connect to a computer running Microsoft Terminal Services. The web site containing the source of the redistributed plug-in is http://properjavardp.sourceforge.net/.
- ssh-plugin.jar—The Secure Shell-Telnet plug-in lets the remote user establish a Secure Shell or Telnet connection to a remote computer. The web site containing the source of the redistributed plug-in is http://javassh.org/.

  Note The ssh-plugin.jar provides support for both SSH and Telnet protocols. The SSH client supports SSH Version 1.0.

- vnc-plugin.jar—The Virtual Network Computing plug-in lets the remote user use a monitor, keyboard, and mouse to view and control a computer with remote desktop sharing turned on. The web site containing the source of the redistributed plug-in is http://www.tightvnc.com.

Note Per the GNU General Public License (GPL), Cisco redistributes plug-ins without having made any changes to them. Per the GPL, Cisco cannot directly enhance these plug-ins.

The security appliance does the following when you install a plug-in onto the flash device:

- (Cisco-distributed plug-ins only) Unpacks the jar file specified in the URL.
- Writes the file to the csco-config/97/plugin directory on the security appliance file system.
- Enables the plug-in for all future clientless SSL VPN sessions, and adds a main menu option and an option to the drop-down menu next to the Address field of the portal page.

When the user in a clientless SSL VPN session clicks the associated menu option on the portal page, the portal page displays a window to the interface and displays a help pane. The user can select the protocol displayed in the drop-down menu and enter the URL in the Address field to establish a connection.

Note Some Java plug-ins might report a status of connected or online even when a session to the destination service is not set up. The open-source plug-in reports the status, not the security appliance.

In the Plug-in tab of the SSL VPN Global Settings page, you can view the currently configured browser plug-ins for clientless SSL VPN browser access. From this tab, you can create or edit the plug-in files, as described in the following procedure.

Plug-in Requirements and Restrictions

Clientless SSL VPN must be enabled on the security appliance to provide remote access to the plug-ins. The minimum access rights required for remote use belong to the guest privilege mode. The plug-ins automatically install or update the Java version required on the remote computer. A stateful failover does not retain sessions established using plug-ins. Users must reconnect following a failover.

Before installing a plug-in, prepare the security appliance as follows:

- Make sure clientless SSL VPN is enabled on an interface on the security appliance.
- Install an SSL certificate onto the security appliance interface to which remote users use a fully-qualified domain name (FQDN) to connect.

Note Do not specify an IP address as the common name (CN) for the SSL certificate. The remote user attempts to use the FQDN to communicate with the security appliance. The remote PC must be able to use DNS or an entry in the System32\drivers\etc\hosts file to resolve the FQDN.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5
- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Plug-in** tab. The Plug-in tab lists all configured plug-ins, including the type of plug-in and the name of the File policy object that defines the actual plug-in file.

Step 3 Do any of the following:

- To add a plug-in, click the **Add Row** button beneath the table and fill in the Add Plug-In Entry dialog box as follows:
  - **Plug-in**—Select the type of plug-in that you are adding:
  - Remote Desktop (RDP) or RDP2—For Remote Desktop Protocol services.
  - Secure Shell (SSH), Telnet—For Secure Shell and Telnet services.
  - VNC—For Virtual Network Computing services.

  - Citrix (ICA)—For Citrix MetaFrame services.
  - Post—For post services.
  - **Plug-in File**—The name of the File policy object that defines the plug-in file. Enter the name of the File object or click **Select** to select an object or to create a new one. For more information on creating File Objects, see Add and Edit File Object Dialog Boxes, page 34-37.

- To edit a plug-in, select it, click the **Edit Row** button, and make your changes in the Edit Plug-In Entry dialog box.
- To delete a plug-in, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Understanding SSL VPN AnyConnect Client Settings

The Cisco AnyConnect VPN Client provides secure SSL and IKEv2 IPsec connections to the security appliance for remote users. The client gives remote users the benefits of an SSL or IKEv2 IPsec VPN client without the need for network administrators to install and configure clients on remote computers.

Tip IKEv2 IPsec connections require AnyConnect 3.0 or higher clients.

Without a previously installed client, remote users enter the IP address in their browser of an interface configured to accept SSL or IKEv2 IPsec VPN connections. Unless the security appliance is configured to redirect http:// requests to https://, users must enter the URL in the form https://<*address*>.

After the user enters the URL, the browser connects to that interface and displays the login screen. If the user satisfies the login authentication, and the security appliance identifies the user as requiring the client, it downloads the client that matches the operating system of the remote computer. After downloading, the client installs and configures itself, establishes a secure connection and either remains or uninstalls itself (depending on the security appliance configuration) when the connection terminates.

In the case of a previously installed client, when the user authenticates, the security appliance examines the revision of the client and upgrades the client as necessary.

When the client negotiates a connection with the security appliance, it connects using Transport Layer Security (TLS), and optionally, Datagram Transport Layer Security (DTLS). DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.

The AnyConnect client can be downloaded from the security appliance, or it can be installed manually on the remote workstation by the system administrator. For more information about installing the client manually, see the *Cisco AnyConnect Secure Mobility Client Administrator Guide*. AnyConnect documentation is available at http://www.cisco.com/en/US/products/ps10884/tsd_products_support_series_home.html. You can find general information about AnyConnect at http://www.cisco.com/go/anyconnect.

The security appliance downloads the client based on the group policy or username attributes of the user establishing the connection. You can configure the security appliance to automatically download the client, or you can configure it to prompt the remote user about whether to download the client. In the latter case, if the user does not respond, you can configure the security appliance to either download the client after a timeout period or present the login page.

AnyConnect Client Profiles

An AnyConnect client profile is a group of configuration parameters, stored in an XML file, that the client uses to configure the connection entries that appear in the client user interface. These parameters (XML tags) include the names and addresses of host computers and settings to enable additional client features.

The AnyConnect client installation includes a profile template, named *AnyConnectProfile.tmpl*, that you can edit with a text editor and use as a basis to create other profile files. You can also set advanced parameters that are not available through the user interface. The installation also includes a complete XML schema file, named *AnyConnectProfile.xsd*.

You can add the profile to the Client Settings tab in the Other Settings policy to have it loaded onto the security appliance and subsequently downloaded to the client workstations based on group policies and username attributes.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Cisco AnyConnect Profile Editor, page 31-63

Cisco AnyConnect Profile Editor

You can configure a profile using the AnyConnect profile editor, a convenient GUI-based configuration tool launched from Cisco Security Manager. The AnyConnect software package for Windows, version 2.5 and later, includes the editor, which activates when you launch the editor from the Add/Edit AnyConnect Client Profile dialog box as long as you have added an appropriate AnyConnect package to the AnyConnect Client Image list.

Note The Cisco AnyConnect Profile Editor is an independent program. For information about configuring AnyConnect profiles, and what AnyConnect Profile Editor can do for you, see the materials available online at http://www.cisco.com/en/US/products/ps10884/products_installation_and_configuration_guides_list.html.

Note Beginning with version 4.7, Security Manager provides support for AnyConnect version 3.2.

Navigation Path

Open the Add/Edit AnyConnect Client Profile dialog box, then click **Launch Editor** (you must first add an appropriate AnyConnect package to the AnyConnect Client Image list before accessing the Add/Edit AnyConnect Client Profile dialog box). The AnyConnect Profile Editor is displayed.

Related Topics

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Understanding and Managing SSL VPN Support Files, page 30-5

Configuring SSL VPN AnyConnect Client Settings (ASA)

This procedure shows you how to define SSL and IKEv2 IPsec VPN client images and profiles. For a detailed explanation of AnyConnect images and profiles, see Understanding SSL VPN AnyConnect Client Settings, page 31-62.

Tip Ensure that you add AnyConnect images of the required releases. For example, if you are configuring an IKEv2 IPsec VPN, you must include an AnyConnect 3.0 or higher image. In general, the image versions must support the features you are deploying in the remote access VPN.

Related Topics

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Cisco AnyConnect Profile Editor, page 31-63
- Understanding and Managing SSL VPN Support Files, page 30-5
- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Client Settings** tab. The tab has two tables listing the configured AnyConnect clients and profiles separately.

The AnyConnect images include an order number. The security appliance downloads portions of the AnyConnect images to the remote computer until it achieves a match with the operating system, starting with the highest order number. Therefore, you should give the highest number to the image used by the most commonly-encountered operating system.

Because mobile users have slower connection speeds, you should load the AnyConnect image for Windows Mobile at the top of the list. Alternatively, you can decrease the connection time by specifying the regular expression **Windows CE** to match the user agent on Windows Mobile devices. When the browser on the mobile device connects to the ASA, it includes the User-Agent string in the HTTP header. The ASA, receiving the string, immediately downloads AnyConnect for Windows Mobile without ascertaining whether the other AnyConnect images are appropriate.

Step 3 To add an AnyConnect client image or make changes to the existing list, do any of the following:

- To add an AnyConnect image, click the **Add Row** button beneath the table and fill in the Add AnyConnect Client Image dialog box. You need to specify the name of the File object that defines the image and the priority order of the image. You can also specify a regular expression for the connecting client to speed up the download. For detailed information about the options, see Add/Edit AnyConnect Client Image Dialog Box, page 31-65.
- To edit an image, select it, click the **Edit Row** button, and make your modifications in the Edit AnyConnect Client Image dialog box.
- To delete an image, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Step 4 To add an AnyConnect profile or make changes to the existing list, do any of the following:

- To add an AnyConnect profile, click the **Add Row** button beneath the table and configure these options in the Add AnyConnect Client Profile dialog box:

- **AnyConnect Profile Name**—The name of the profile.

  To use this profile, ensure that you specify the profile name in an ASA Group Policy object assigned to the security appliance (in the Full Client settings page as described in ASA Group Policies SSL VPN Full Client Settings, page 34-19). Configure the ASA Group Policy object through the remote access Connection Profiles policy for the device as described in Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.

- **AnyConnect Profile Type**—Select the type of AnyConnect Profile you are adding or editing: VPN, Network Access Manager, Telemetry, Web Security, ISE Posture, or Customer Experience Feedback.

- **AnyConnect Profile File**—The name of the File object that identifies the Anyconnect client profile XML file. Click **Select** to select the object or to create a new one. For more information about File objects, see Add and Edit File Object Dialog Boxes, page 34-37.

  Note Beginning with version 4.7, Security Manager provides support for AnyConnect version 3.2. If you select ISE Posture as the AnyConnect Profile Type, the AnyConnect Profile File will have a filename extension of .isp.

  Note If you are going to create a new profile using the AnyConnect Profile Editor, do not specify an AnyConnect Profile File.

- **Enable Storage URL**—Beginning with version 4.12, Security Manager enables you to select either Private or Shared option for ASA 9.6(2) or later Multi-Context devices.

- **Launch Editor**—Click **Launch Editor** to use the AnyConnect Profile Editor to edit the profile specified in AnyConnect Profile File or to create a new profile if no profile file is specified. For more information about File objects, see Cisco AnyConnect Profile Editor, page 31-63.

- To edit a profile, select it, click the **Edit Row** button, and make your modifications in the Edit AnyConnect Client Profile dialog box.
- To delete a profile, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Add/Edit AnyConnect Client Image Dialog Box

Use the Add or Edit AnyConnect Client Image dialog box to create or edit a package file as the client image, and establish the order that the security appliance downloads the image to the remote workstation.

Navigation Path

From the Client Settings tab of the SSL VPN Other Settings policy for ASA devices, click the **Add Row** button for the AnyConnect Client Image table, or select an image and click the **Edit Row** button. For detailed information on opening the tab, see Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64.

Related Topics

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Understanding and Managing SSL VPN Support Files, page 30-5

Field Reference

Table 31-23 Add or Edit AnyConnect Client Image Dialog Box

Element	Description
AnyConnect Client Image	The name of the File object that identifies the Anyconnect client. Click **Select** to select an object or to create a new one. For more information about File objects, see Add and Edit File Object Dialog Boxes, page 34-37.
Image Order	The order in which the security appliance downloads the client images to the remote workstation. It downloads the image in priority order. Therefore, you should enter a lower value for the image used by the most commonly-encountered operating system.
Regular Expression	A regular expression to match the user agent. Enter a name of an existing regular expression policy object or click **Select** to select an entry from the Regular Expressions Selector dialog box. To add a new regular expression, click the **Add (+)** button on the Regular Expressions Selector dialog box. For more information see Add/Edit Regular Expressions, page 17-108. If you are adding an AnyConnect package for Windows Mobile, specify the regular expression **Windows CE** to match the user agent on Windows Mobile devices. This decreases the connection time of the mobile device. When the browser on the mobile device connects to the adaptive security appliance, it includes the User-Agent string in the HTTP header. The adaptive security appliance, receiving the string, immediately downloads AnyConnect for Windows Mobile without ascertaining whether the other AnyConnect images are appropriate.
Enable Storage URL (Only for ASA 9.6(2) or later Multi-Context devices)	Beginning with version 4.12, Security Manager enables you to select either Private or Shared option for ASA 9.6(2) or later Multi-Context devices.

Understanding Kerberos Constrained Delegation (KCD) for SSL VPN (ASA)

There are many ways to protect network resources through the use of authentication. Many organizations want to use Kerberos to protect certain web applications while using other authentication techniques, such as username and password, digital certificates, RSA SecureID, or SmartCards, to control access to an SSL VPN. However, a restriction in the Kerberos protocol prevents Kerberos authentication if the user has already used another technique to authenticate to the VPN.

Microsoft overcomes this limitation in Kerberos starting with Windows Server 2003. Using protocol transition and constrained delegation, the ASA can authenticate to the Kerberos Key Distribution Center (KCD) on the Windows domain controller and obtain impersonate tickets for users who have authenticated to the ASA using non-Kerberos protocols. The ASA can use the impersonate ticket to obtain other Kerberos service tickets for remote users.

To configure the domain controller so that Kerberos constrained delegation works, you must do the following:

- Each instance of a service that uses Kerberos authentication must have a service principle name (SPN) defined so that clients can identify it on the network. Register the SPN in the Active Directory **Service-Principal-Name** attribute of the Windows account under which the instance of the service is running. When a service needs to authenticate to another service running on a specific computer, it uses that service's SPN to differentiate it from other services running on that computer.

  The SPN syntax is *service_class/host_name:port*, where:

  - *service_class* identifies the service. It can be a built-in service, such as http, or a user defined service.
  - *host_name* identifies the fully-qualified domain name or NetBIOS name of the server that hosts the service, but it cannot be an IP address.
  - *port* identifies the port on which the service runs. You can omit the port if you use the default service port.
- Create a service account username and password that the ASA can use. Configure the account to allow Kerberos constrained delegation to any authentication protocol. In addition, the user account must not be marked as a sensitive account that cannot be delegated.

  To configure the ASA to allow KCD, once the ASA joins the domain, an entry should appear under the Users and Computers list on the domain controller for the ASA. In the Properties dialog box, on the Delegation tab, select **Trust this computer for delegation to specified services only**, and then select **Use any authentication protocol**. In the table of authorized services, add all services for which the ASA is delegated for authentication on behalf of users.

Tip For definitive information on configuring this feature on the Windows domain controller, refer to the Microsoft documentation.

For the ASA to use Kerberos constrained delegation, you must configure the ASA as described in Configuring Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-69. The feature is available on ASA Software release 8.4 and higher only.

Following is an example that explains how Kerberos constrained delegation works with a clientless SSL VPN hosted on an ASA.

Figure 31-1 Kerberos Constrained Delegation Example

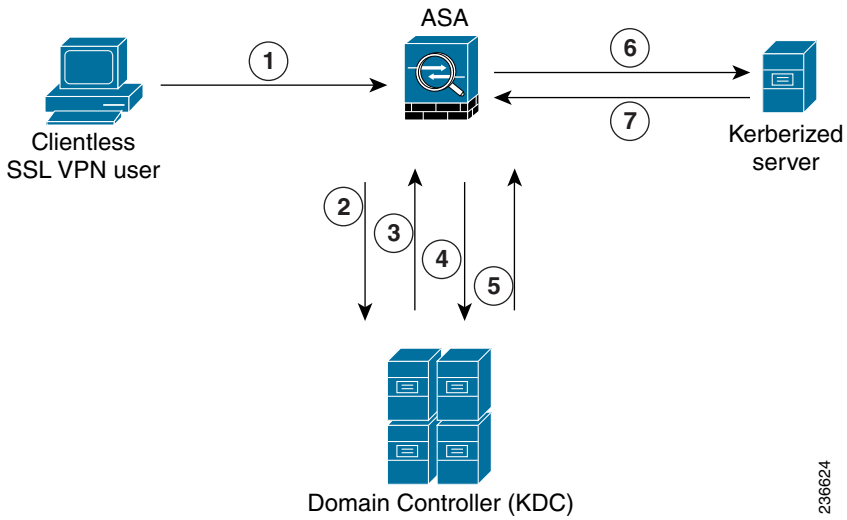

After verifying the identity of an SSL VPN user with the configured authentication mechanism, the ASA uses protocol transition to switch to the Kerberos protocol for authentication on behalf of the user and then sends a Kerberos service ticket instead of the user's credentials to a published Web server that accepts Kerberos for authentication. Following are the steps:

1. An SSL VPN user session is authenticated by the ASA using the authentication mechanism configured for the user. For example, in case of Smartcard credentials, the ASA extracts the required information (the user's principle name) from the digital certificate and performs LDAP authorization against Windows Active Directory.
2. After successful authentication, the user logs into the ASA SSL VPN portal page. The VPN user accesses a web service by entering a URL in the portal page or by clicking on a bookmark. If the access requires authentication, the server challenges the ASA for credentials and along with the challenge sends a list of authentication mechanisms supported by the server. Based on the HTTP headers in the challenge, the ASA deduces whether the server requires Kerberos authentication. If connecting to a backend server requires Kerberos authentication, then the ASA requests an impersonate ticket, for itself on behalf the user, from the KDC.
3. The KDC returns the requested tickets to the ASA. Even though these tickets are passed to the ASA, they contain the user's authorization data.

   Note These first steps comprise protocol transition; after these steps, a user who authenticated to the ASA using a non-Kerberos authentication protocol is transparently authenticated to the KDC using Kerberos.

4. The ASA now requests a service ticket from the KDC for the specific service that the user wants to access. The service ticket request contains the SPN (the unique identifier) of the service.
5. The KDC returns a service ticket for the specific service to the ASA.
6. The ASA uses the service ticket to request access to the web service, in the above scenario this is sent to the web server in a HTTP GET request.
7. The web server authenticates the Kerberos service ticket and grants access to the service. An authentication failure will display an appropriate error message after acknowledgment of which the portal will be displayed.

Configuring Kerberos Constrained Delegation (KCD) for SSL VPN (ASA)

Use the Microsoft KCD Server tab of the SSL VPN Other Settings page to configure Kerberos Constrained Delegation (KCD) for clientless SSL VPNs hosted on an ASA.

KCD addresses a limitation of Kerberos. If a user authenticates to the SSL VPN using a method other than Kerberos, the user cannot access Kerberos-protected resources. This prevents a remote access device, such as ASA, from authenticating users using non- Kerberos methods and still provide single sign-on access to Kerberos-authenticated web applications in the enterprise.

If this limitation applies to your network, you can configure KCD to get around the limitation. KCD offloads the Kerberos authentication to the ASA. Users log into the corporate network using the SSL VPN portal, and from then on, access Kerberos-protected services in a transparent fashion.

Tips

- KCD requires ASA release 8.4+. If you configure KCD for other releases, the configuration is ignored.
- The feature is used with clientless SSL VPN access only.
- Microsoft Windows Server (2003 or 2008), configured as domain controllers, are required for KCD.
- If you use SSL VPN Bookmark policy objects to define bookmarks to include on the SSL VPN portal page, you might need to add explicit service principle name (SPN) parameters to bookmarks if a service uses a non-default port. For services that use Kerberos authentication, an SPN must be defined in the Service-Principle-Name attribute of the account under which the service runs.

  Bookmarks need to reflect this configuration. The SPN is a parameter on the URL: http://<url>?SPN=<spn> or http://<url>?SPN=<spn>. For example, **http://owa.example.com?SPN=http/owa:444**. For more details about the SPN syntax, see Understanding Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-66.
- To configure this feature, you must also configure the Hostname, DNS, and NTP policies. Configure both hostname and domain name in the Hostname policy.
- Kerberos authentication requires that the clock between the hosts to be synchronized with a maximum drift of 5 minutes (this is the default setting). This restriction is applicable to the clocks on the ASA, the domain controller, and the application servers. Configuring the same NTP server for all servers should address the requirement.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51
- Understanding AAA Server and Server Group Objects, page 6-27

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Microsoft KCD Server** tab.

Step 3 Select **Configure KCD** and configure the following options:

- **KCD Server**—The AAA server group object that identifies the Microsoft KCD server (the domain controller) to use for Kerberos Constrained Delegation. Enter the name of the object or click **Select** to select it from a list or to create a new object. The object must use a Kerberos AAA server policy object to identify the domain controller.
- **Username, Password, Confirm**—A user account that the ASA can use to join the Active Directory domain.

  For the ASA to use Kerberos protocol transition and constrained delegation, and obtain service tickets on behalf of the remote access users, the account used by the ASA to authenticate with the domain controller must be configured in Active Directory and configured to allow Kerberos constrained delegation to any authentication protocol. In addition, the user account must not be marked as a sensitive account that cannot be delegated. For more information about Active Directory configuration requirements, see Understanding Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-66.

Configuring AnyConnect Custom Attributes (ASA)

AnyConnect custom attributes allow for a more expeditious delivery and deployment of new endpoint features by giving the ASA the ability to generically support the addition of new client controls without the need for an ASA software upgrade.

In the AnyConnect Custom Attribute tab of the SSL VPN Other Settings page, you can view configured AnyConnect custom attributes, add new attributes, and modify or delete existing attributes.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5
- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **AnyConnect Custom Attribute** tab. The AnyConnect Custom Attribute tab lists all defined custom attributes.

Step 3 Do any of the following:

- To add a custom attribute, click the **Add Row** button beneath the table and fill in the Add AnyConnect Custom Attribute dialog box. The options are described in detail in Add/Edit AnyConnect Custom Attribute Dialog Box, page 31-71.
- To edit a custom attribute, select it, click the **Edit Row** button, and make your changes in the Edit Plug-In Entry dialog box.
- To delete a custom attribute, select it and click the **Delete Row** button. You are asked to confirm the deletion.

Add/Edit AnyConnect Custom Attribute Dialog Box

Use the Add or Edit AnyConnect Custom Attribute dialog box to add or modify an AnyConnect custom attribute. AnyConnect custom attributes allow for a more expeditious delivery and deployment of new endpoint features by giving the ASA the ability to generically support the addition of new client controls without the need for an ASA software upgrade.

Beginning with version 4.7, Security Manager enables to add Custom Attribute Data to an existing Custom Attribute Type for ASA devices running the software version 9.3(1) or higher. Use the Add or Edit AnyConnect Custom Attribute Data dialog box to add or modify the attribute name and attribute value for an existing AnyConnect custom attribute type. For more information see Add/Edit AnyConnect Custom Attribute Data Dialog Box, page 31-71.

Navigation Path

From the AnyConnect Custom Attribute tab of the SSL VPN Other Settings policy for ASA devices, click the **Add Row** button for the AnyConnect Custom Attributes table, or select an attribute and click the **Edit Row** button. For detailed information on opening the tab, see Configuring AnyConnect Custom Attributes (ASA), page 31-70.

Related Topics

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Understanding and Managing SSL VPN Support Files, page 30-5

Field Reference

Table 31-24 Add or Edit AnyConnect Custom Attributes Dialog Box

Element	Description
Type	The type of the AnyConnect custom attribute. This is used when referencing the attribute in Security Manager and in the aggregate auth protocol messages sent to the AnyConnect client. The maximum length is 32 characters.
Description	A free form description of attribute usage. This text will appear in the command help when the custom attribute is referenced from the group-policy attribute configuration mode. The maximum length is 128 characters.

Add/Edit AnyConnect Custom Attribute Data Dialog Box

Beginning with Security Manager version 4.7, you can use the Add or Edit AnyConnect Custom Attribute Data dialog box to add or modify the attribute name and attribute value for an existing AnyConnect custom attribute type.

Navigation Path

Click the AnyConnect Custom Attribute tab of the SSL VPN Other Settings policy for ASA devices. On the Custom Attribute table, select an attribute type and then click the **Add Row** button for the Custom Attribute Data table. Or on the Custom Attribute Data table select an existing Custom Attribute Data and click the **Edit Row** button.

For each attribute type, you can define multiple attribute names with corresponding values.

For information on adding or modifying an attribute type, see Understanding SSL VPN AnyConnect Client Settings, page 31-62.

Related Topics

- Understanding SSL VPN AnyConnect Client Settings, page 31-62
- Add/Edit AnyConnect Custom Attribute Dialog Box, page 31-71
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Understanding and Managing SSL VPN Support Files, page 30-5

Field Reference

Table 31-25 Add or Edit AnyConnect Custom Attribute Data Dialog Box

Element	Description
Attribute Name	The name of an AnyConnect custom attribute. The name is used when referencing the attribute in group-policy and dynamic-access-policy-record config mode. The maximum length is 32 characters.
Attribute Value	A free form string containing the attribute value. The attribute value is associated with the attribute name and it is passed to the client during the configuration of the connection. The maximum length of the string can be 420 characters. The attribute value can contain multiple text lines.

Configuring SSL VPN Advanced Settings (ASA)

Use the Advanced tab of the SSL VPN Other Settings page to configure the memory, on-screen keyboard, and internal password features on ASA devices. Beginning with Cisco Security Manager 4.15, you can also enable the HSTS support and specify the timeout values. All of these settings are optional.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 On the Other Settings page, click the **Advanced** Tab.

Step 3 In the **Memory Size** field, specify the amount of memory you want to allocate to SSL VPN sessions. The default is 50%.

To change the setting, select one of the following options and enter the desired number:

- **% of Total Physical Memory**—As a percentage of total memory. Default is 50%.
- **Kilobytes**—In kilobytes. 20KB is the minimum setting allowed. Cisco recommends that you do not specify memory in terms of KB because different ASA models have different total amounts of memory, for example:

Note When you change the memory size, the new setting takes effect only after the system reboots.

Step 4 In the **Enable On-Screen Keyboard** field, select one of the following options:

- **Disabled**—The on-screen keyboard is not displayed. Users must input their credentials using the standard keyboard. This is the default.
- **On All Pages**—Allows a user to input credentials using an on-screen keyboard, which is displayed whenever logon credentials are required.
- **On Logon Page Only**—Allows a user to input credentials using an on-screen keyboard, which is displayed on the logon page but not on any other pages that require credentials.

Step 5 Select **Allow Users to Enter Internal Password** to require an additional password when accessing internal sites. This feature is useful if you require that the internal password be different from the SSL VPN password. For example, you can use a one-time password for authentication to ASA and another password for internal sites.

Note The HSTS option is available for ASA 9.8.2 devices only.

Step 6 In the HTTP Strict Transport Security (HSTS) area, do the following:

- To enable HSTS support, select the **Enable HSTS Header** check box. HSTS feature can be enabled by sending header to clients. To disable the support, clear this check box.
- If you have selected the **Enable HSTS Header** check box, in **HSTS Header Timeout** enter the timeout value. If you leave this field blank, Cisco Security Manager will use the default Timeout value of 10886400.

Configuring SSL VPN Server Verification (ASA)

When connecting to a remote SSL-enabled server through clientless SSL VPN, it is important to know that you can trust the remote server, and that it is in fact the server you are trying to connect to. ASA 9.0 introduces support for SSL server certificate verification against a list of trusted certificate authority (CA) certificates for clientless SSL VPN.

When you connect to a remote server via a web browser using the HTTPS protocol, the server will provide a digital certificate signed by a CA to identify itself. Web browsers ship with a collection of CA certificates which are used to verify the validity of the server certificate. This is a form of public key infrastructure (PKI).

Just as browsers provide certificate management facilities, so does the ASA in the form of trusted certificate pool management facility: trustpools. This can be thought of as a special case of trustpoint representing multiple known CA certificates. The ASA includes a default bundle of certificates, similar to that provided with web browsers, but it is inactive until activated by the administrator.

Note If you are already familiar with trustpools from Cisco IOS then you should be aware that the ASA version is similar, but not identical.

This procedure describes how to enable HTTPS server verification for clientless SSL VPN users.

Related Topics

- Configuring Other SSL VPN Settings (ASA), page 31-51
- Configuring Trusted Pool Settings (ASA), page 31-30
- Using the Trustpool Manager, page 31-32

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Other Settings** from the Policy selector. Click the **SSL Server Verification** tab.
- (Policy view) Select **Remote Access VPN > SSL VPN > Other Settings (ASA)** from the Policy Type selector. Select an existing policy or create a new one. Click the **SSL Server Verification** tab.

Step 2 Select **Enable** to enable HTTPS Server Verification for Clientless SSL VPN users.

Step 3 Specify the action you want to be taken if server certificate verification fails:

- **Disconnect user from Https page** – Disconnect if the server could not be verified.
- **Allow user to continue to Https page** – Allow the user to continue the connection, even if the check failed.

Configuring SSL VPN Shared Licenses (ASA 8.2+)

Use the SSL VPN Shared License page to configure your SSL VPN Shared License.

You can purchase a shared license with a large number of SSL or remote access IKEv2 IPsec VPN sessions and share the sessions as needed among a group of ASA devices by configuring one of the ASA devices as a shared license server, and the rest as clients. For the server license, you can share 500-50,000 licenses in increments of 500 and 50,000-1,040,000 licenses in increments of 1000.

A license is consumed by each remote access user that makes an SSL or IKEv2 IPsec connection.

Note The shared license cannot be used at the same time as the AnyConnect Essentials license.

The following topics explain the procedure for configuring shared licenses:

- Configuring an ASA Device as a Shared License Client, page 31-76
- Configuring an ASA Device as a Shared License Server, page 31-77

Navigation Path

- (Device View) Select an ASA device using version 8.2 or higher, and select **Remote Access VPN > SSL VPN > Shared License** from the Policy selector.
- (Policy View) Select **Remote Access VPN > SSL VPN > Shared License (ASA 8.2+)** from the Policy Type selector. Select an existing policy or create a new one.

Field Reference

Table 31-26 SSL VPN Shared License Page

Element	Description
Select Role	The role you are configuring, either Shared License Client or Shared License Server. Depending on your choice, different fields appear.
Shared License Client	
Shared Secret	The case-sensitive string (4-128 characters) used for communicating with the shared license server.
License Server	The IP address or the name of a network/host object that identifies the ASA device configured as the license server. Click **Select** to select an existing object or to create a new one.
License Server Port	The number of the TCP port on which the license server communicates. Enter a port number or the name of a port list object, or click **Select** to select an object or to create a new one.
Select Backup Role of Client	The backup role of the client: • Client Only—When selected, the client acts only as the client. In this case, you can specify another device as a backup server. Enter the IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new object. • Backup Server—When selected, the client also acts as the backup server. In this case, you must also specify the interfaces to be used for this purpose. Enter a comma-separated list of interface names or interface role objects, or click **Select** to select interfaces or objects or to create new objects.
Shared License Server	
Shared Secret	The case-sensitive string (4-128 characters) used for communicating with the shared license server.
License Server Port	The number of the TCP port on which the license server communicates. Enter a port number or the name of a port list object, or click **Select** to select an object or to create a new one.
Refresh Interval	The refresh interval, between 10-300 seconds. The default is 30 seconds.
Interfaces	A comma-separated list of interfaces used for communicating shared licenses to clients. Enter the names of interfaces or interface role objects, or click **Select** to select interfaces or objects or to create new objects.

Table 31-26 SSL VPN Shared License Page (continued)

Element	Description
Configure Backup shared SSL VPN License Server	Whether to configure a backup server for the shared license server. If you select this option, configure the following: • **Backup License Server**—The IP address, or network/host object that contains the address, of the server to act as a backup license server if the current one is unavailable. Click **Select** to select an object or to create a new one. • **Backup Server Serial Number**—The serial number of the backup license server. • **HA Peer Serial Number**—(Optional) The serial number of the backup server of a failover pair.

This section contains the following topics:

- Configuring an ASA Device as a Shared License Client, page 31-76
- Configuring an ASA Device as a Shared License Server, page 31-77

Configuring an ASA Device as a Shared License Client

This procedures describes how to configure an ASA device as a shared license client.

Tip You must ensure that the SSL VPN Shared License Client activation key is present on the device.

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Shared License** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN >Shared License (ASA 8.2+)** from the Policy Type selector. Select an existing policy or create a new one.

The SSL VPN Shared License page appears (see Configuring SSL VPN Shared Licenses (ASA 8.2+), page 31-74).

Step 2 Select **Shared License Client** as the role of the device.

Step 3 In the Shared Secret field, enter and confirm a case-sensitive string (4-128 characters) used for communicating with the shared license server.

Step 4 In the License Server field, enter the IP address or the name of a network/host object that identifies the ASA device configured as the license server.

Step 5 In the License Server Port field, enter the number of the TCP port on which the license server communicates.

Step 6 Select the role of the client:

- **Client Only**—When selected, the client acts only as the client. In this case, you can specify another device as a backup server.

- **Backup Server**—When selected, the client also acts as the backup server. In this case, you must also specify the interfaces to be used for this purpose.

Configuring an ASA Device as a Shared License Server

This procedures describe how to configure an ASA device as a shared license server.

Tip You must ensure that the SSL VPN Shared License Server activation key is present on the device.

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > SSL VPN > Shared License** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN >Shared License (ASA 8.2+)** from the Policy Type selector. Select an existing policy or create a new one.

The SSL VPN Shared License page appears (see Configuring SSL VPN Shared Licenses (ASA 8.2+), page 31-74).

Step 2 Select **Shared License Server** as the role of the device.

Step 3 In the Shared Secret field, enter and confirm a case-sensitive string (4-128 characters) used for communicating with the shared license server.

Step 4 In the License Server Port field, enter the number of the TCP port on which the license server communicates.

Step 5 In the Refresh Interval field, enter a value between 10-300 seconds to be used as the refresh interval. Default is 30 seconds.

Step 6 In the Interfaces field, enter or select the interfaces to be used for communicating with clients.

Step 7 (Optional.) Select **Configure Backup shared SSL VPN License Server** to configure a backup server for the shared license server, then configure the following:

- **Backup License Server**—The IP address, or network/host object that contains the address, of the server to act as a backup license server if the current one is unavailable.
- **Backup Server Serial Number**—The serial number of the backup license server.
- **HA Peer Serial Number**—(Optional) The serial number of the backup server of a failover pair.

Customizing Clientless SSL VPN Portals

You can customize the web site and its contents that you use for the portal page for a browser-based clientless SSL VPN. ASA devices allow much more customization than IOS devices. You can create several policy objects that define the look of the web pages the user sees when logging into or out of the VPN and the home page for the portal, as well as the bookmarks and smart tunnels available to the user.

This section contains the following topics:

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

- Localizing SSL VPN Web Pages for ASA Devices, page 31-80
- Creating Your Own SSL VPN Logon Page for ASA Devices, page 31-82
- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs, page 31-88

Configuring ASA Portal Appearance Using SSL VPN Customization Objects

An SSL VPN Customization object describes the appearance of browser-based clientless SSL VPN web pages displayed to users. This includes the Logon page displayed when they connect to the ASA security appliance, the Home page displayed after authentication, and the Logout page displayed when users log out of the SSL VPN service.

You use SSL VPN Customization objects when defining ASA group objects or Remote Access VPN Connection policies for ASA devices. You can create several customization objects and define multiple ASA group or connection profiles so that each user group sees web pages designed specifically for their use. Customization can include localizing the web pages in the languages appropriate for each group. For more information about localization, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

Initially, when a user first connects, the default customization object identified in the connection profile determines how the logon screen appears. If the user selects a different group from the connection profile list on the logon page, and that group has its own customization, the screen changes to reflect the customization object for the selected group. After the remote user is authenticated, the screen appearance is determined by the customization object that has been assigned to the group policy.

After you create the SSL VPN customization object as described in this procedure, you can use the object to specify the portal characteristics in these policies:

- On the **SSL VPN > Settings** page in an ASA group policy object (see ASA Group Policies SSL VPN Settings, page 34-25), which you then select in one of these policies:
  - **Remote Access VPN > Group Policies**
  - **Remote Access VPN > Connection Profiles** on the **General** tab
- In the **Remote Access VPN > Connection Profiles** policy, you can also specify the SSL VPN customization object on the **SSL** tab (see SSL Tab (Connection Profiles), page 31-22).

Related Topics

- Localizing SSL VPN Web Pages for ASA Devices, page 31-80
- Creating Policy Objects, page 6-9
- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Tip You can also create SSL VPN Customization objects when defining policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Step 2 Select **SSL VPN Customization** from the Object Type selector. The SSL VPN Customization page opens, displaying a list of the existing SSL VPN Customization objects.

Step 3 Right-click in the work area and select **New Object**.

The Add SSL VPN Customization dialog box appears (see Add and Edit SSL VPN Customization Dialog Boxes, page 34-51).

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Before you configure settings for the various pages, use the Preview button to view the default settings. Clicking **Preview** opens a browser window to display the current settings for the Logon page, Portal page, or Logout page, whichever one is selected in the table of contents (selecting a page within one of these folders is the same as selecting the parent folder).

Tip Click **Preview** after making any changes to settings to verify that the changes are what you desire.

Step 6 Configure the settings for the Logon page. This web page is the one users see first when connecting to the SSL VPN portal. It is used for logging into the VPN. Select the following items in the Logon Page folder in the table of contents on the left of the dialog box to view and change the settings:

- **Logon Page**—Specify the title of the web page, which is displayed in the browser's title bar.
- **Title Panel**—Determine whether the Logon page will have a title displayed in the web page itself. If you enable the title panel, you can specify the title, font, font size and weight, styles, and colors used. You can also select a File object that identifies a logo graphic. For more information about the settings, see SSL VPN Customization Dialog Box—Title Panel, page 34-53.
- **Language**—If you want to configure translation tables for other languages on the ASA device and use them, you can configure the supported languages and allow users to choose their language. For information about translation tables and localization support, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80. For more information about the settings, see SSL VPN Customization Dialog Box—Language, page 34-54.
- **Logon Form**—Configure the labels and colors used in the form that accepts user logon information. For more information about the settings, see SSL VPN Customization Dialog Box—Logon Form, page 34-56.
- **Informational Panel**—If you want to provide extra information to the user, you can enable an informational panel and add text and a logo graphic. For more information about the settings, see SSL VPN Customization Dialog Box—Informational Panel, page 34-57.
- **Copyright Panel**—If you want to include copyright information on the logon page, you can enable the copyright panel and enter your copyright statement. For more information about the settings, see SSL VPN Customization Dialog Box—Copyright Panel, page 34-58.
- **Full Customization**—If you do not want to use the security appliance's built-in logon page, even customized, you can instead enable full customization and supply your own web page. For information on creating the required file, see Creating Your Own SSL VPN Logon Page for ASA Devices, page 31-82. For more information about the settings, see SSL VPN Customization Dialog Box—Full Customization, page 34-59.

Step 7 Configure the settings for the Portal page. This is the home page for the SSL VPN portal, and is displayed after the users log in. Select the following items in the Portal Page folder in the table of contents on the left of the dialog box to view and change the settings:

- **Portal Page**—Specify the title of the web page, which is displayed in the browser's title bar.

- **Title Panel**—Determine whether the Portal page will have a title displayed in the web page itself. If you enable the title panel, you can specify the title, font, font size and weight, styles, and colors used. You can also select a File object that identifies a logo graphic. For more information about the settings, see SSL VPN Customization Dialog Box—Title Panel, page 34-53.
- **Toolbar**—Determine whether the Portal page will have a toolbar, which contains a field for entering a URL to browse. For more information about the settings, see SSL VPN Customization Dialog Box—Toolbar, page 34-59.
- **Applications**—Determine which application buttons will appear on the page. For more information about the settings, see SSL VPN Customization Dialog Box—Applications, page 34-60.
- **Custom Panes**—Determine how you want to organize the body of the Portal page. The default is a single column with no internal panes. You can create a multiple-column layout, create internal panes that display text or references to URLs, and determine in which column and row to place the panes. For more information about the settings, see SSL VPN Customization Dialog Box—Custom Panes, page 34-61.
- **Home Page**—Determine how and whether to display URL lists on the home page, and whether to use your own web page for the main body of the Portal page. For more information about the settings, see SSL VPN Customization Dialog Box—Home Page, page 34-62.

Step 8 Select **Logout Page** to configure the settings of the page displayed when a user logs out of the SSL VPN. You can configure the title, message text, fonts, and colors. For more information about the settings, see SSL VPN Customization Dialog Box—Logout Page, page 34-63.

Step 9 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 10 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 11 Click **OK** to save the object.

Localizing SSL VPN Web Pages for ASA Devices

Localization is the process of providing text in a language that is appropriate for the target users. When you create an SSL VPN Customization object for defining the look of browser-based clientless SSL VPN web pages hosted on an ASA device, you can configure the pages to use the desired language.

To see localized web pages correctly, users must configure their browsers to use UTF-8 encoding (for example, in Internet Explorer, select **View > Encoding > Unicode (UTF-8)**. They also must install the required fonts or language support files for their language using the Regional and Language Options control panel. On the Languages tab, click Details to install the desired languages, and select the appropriate supplemental language settings for East Asian, complex scripting, and right-left languages. On the Advanced tab, select the desired code page conversion tables. If users do not configure the browser correctly, they might see boxes instead of characters.

There are two techniques you can use to localize SSL VPN web pages that are hosted on an ASA device. These techniques are not mutually exclusive; you can use both of them. These are the techniques:

- **Configure the SSL VPN Customization object using the desired language**—When you create the SSL VPN Customization object, you can enter text for labels and messages in non-English, non-ASCII languages in UTF-8 encoding. To enter non-ASCII languages in UTF-8 encoding, you must configure Windows with the correct locale setting and have the required fonts installed. Use the Regional and Language Options control panel to configure your system and install files required

for complex script or East Asian languages. If you want to type in text directly, you also need to install an appropriate keyboard; otherwise, you can use a text editor that supports the language's characters and copy and paste text from a document that contains the text you want to use.

You can also enter non-ASCII languages into SSL VPN Bookmarks objects.

- **Configure translation tables on the ASA device to support the languages you want to make available**—To enable the security appliance to provide language translation for the portal and screens displayed to users, you must define the necessary languages in a translation table and import the table into the security appliance. The software image package for the security appliance includes a translation table template. Every language you list in an SSL VPN Customization object must have a corresponding translation table on the device. Conversely, translation tables for languages that are not listed in the SSL VPN Customization object are ignored.

  If you use this technique, you must use the ASA CLI or ASDM to configure and upload the translation tables. You cannot manage the translation tables with Security Manager. However, the SSL VPN Customization object includes settings that allow you to configure automatic browser language selection and to enable users to select their desired language. Thus, if you install translation tables for ten languages, the pages defined in the SSL VPN Customization object will be available to users in all of those languages. For more information on these settings, see SSL VPN Customization Dialog Box—Language, page 34-54.

  Although both of the following features require translation tables, they are separate and complementary:

  – **Automatic Browser Language Selection**—Automatic browser language selection attempts to select the appropriate language based on the user's browser settings. This technique does not ask for user input. In the SSL VPN Customization object, you create a list of languages that will be used in the negotiation with the browser. During a connection, the security appliance receives a list of languages (and their priorities) from the browser, and looks through your list of languages top to bottom until a match is found. If there is no match, then the language you defined in the list as the default language is used. If you do not specify a default language, English is used.

    The languages on the security appliance are labels for the translation tables. The languages must mirror those of the browser, and can consist of groups of up to 8 alphanumeric characters (starting from alpha characters), separated by hyphens. For example, fr-FR-paris-univ8. However, when you add a language to the list in Security Manager, only the first two characters are available.

    When looking for a match, the security appliance starts with the longest language name, and if it does not match, it discards the rightmost group of the name. For example, if the preferred language on the browser is fr-FR-paris-univ8, and the security appliance supports fr-FR-paris-univ8, fr-FR-paris, fr-FR, and fr, it matches fr-FR-paris-univ8 and uses the translated strings from that translation table. If fr is the only language on the security appliance, the security appliance considers it a match also, and uses that translation table.

    For more information about setting up translation tables, see the user documentation for the ASA device and operating system or the ASDM online help.

  – **Language Selector**—By enabling the language selector, you provide the user with the ability to actively select the desired language from a list of languages that you support. This technique does not rely on the browser language settings being configured correctly. The language selector is displayed on the logon page.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

- Creating Policy Objects, page 6-9
- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51

Creating Your Own SSL VPN Logon Page for ASA Devices

You can create your own custom SSL VPN Logon page rather than use the page provided by the security appliance for browser-based clientless SSL VPNs. This is called full customization, and replaces the settings you can configure in the SSL VPN Customization policy object.

To provide your own Logon page, you must create the page, copy it to the Security Manager server, and identify the page on the Full Customization page of the SSL VPN Customization object dialog box. For information on creating SSL VPN Customization objects, see Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78.

When you enable full customization, all other settings for the Logon page configured in the policy object are ignored. When you deploy your configuration to the ASA device, Security Manager copies your custom page to the device.

The Logon page you create must include all of the HTML code required to present the page correctly, and include special Cisco HTML code that provides the functions for the login form and the Language Selector drop-down list. Keep the following in mind when you create the HTML file:

- The file extension must be **.inc**.
- All images in the custom Logon page must be present on the security appliance. Replace the file path with the keyword **/+CSCOU+/**, which is an internal directory on the ASA device. When you upload an image to the device, it is saved in this directory.
- Use the **csco_ShowLoginForm('lform')** Javascript function to add the login form to the page. This form prompts for the username, passwords, and group information. You must include this function somewhere on the page.
- Use the **csco_ShowLanguageSelector('selector')** Javascript function to add the Language Selector drop-down list to the page. You do not have to use this function if you are not supporting the use of more than one language.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78
- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51
- SSL VPN Customization Dialog Box—Full Customization, page 34-59

Configuring SSL VPN Bookmark Lists for ASA and IOS Devices

When you configure a browser-based clientless SSL VPN, you can define a list of bookmarks, or URLS, to include on the SSL VPN portal page. Use SSL VPN bookmarks policy objects to define bookmark lists.

You can create SSL VPN bookmark objects for SSL VPNs hosted on IOS devices or ASA devices. However, these device types allow different bookmark configurations, the ASA device allowing more configuration options than IOS devices. Besides allowing more configuration options, you can also create bookmarks for ASA devices in non-English, non-ASCII languages. For more information on localizing the bookmarks and portal for ASA devices, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

After you create the SSL VPN bookmark object as described in this procedure, you can use the object to specify the bookmark object in the **Portal Web Pages** or **Bookmarks** fields in these policies:

- ASA devices—On the **SSL VPN > Clientless** page in an ASA group policy object (see ASA Group Policies SSL VPN Clientless Settings, page 34-12), which you then select in one of these policies:
  - **Remote Access VPN > Group Policies**
  - **Remote Access VPN > Connection Profiles** on the **General** tab
- ASA devices—In the **Remote Access VPN > Dynamic Access** policy, you can specify the SSL VPN bookmark object on the **Main > Bookmarks** tab (see Main Tab, page 32-14).
- IOS devices—On the **Clientless** page in a user group policy object configured for SSL VPN (see User Group Dialog Box—Clientless Settings, page 34-83), which you then select in the **Remote Access VPN > SSL VPN** policy on the **General** tab.

Related Topics

- Creating Group Policies (ASA, PIX 7.0+), page 31-28
- Configuring Dynamic Access Policies, page 32-2
- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Configuring an SSL VPN Policy (IOS), page 33-14
- Creating Policy Objects, page 6-9
- Policy Object Manager, page 6-4

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Tip You can also create SSL VPN bookmark objects when you define policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Step 2 Select **SSL VPN Bookmarks** from the Object Type selector. The SSL VPN Bookmarks page opens, displaying a list of the existing SSL VPN bookmark objects.

Step 3 Right-click in the work area, then select **New Object**.

The Add SSL VPN Bookmark dialog box appears (see Add or Edit Bookmarks Dialog Boxes, page 34-44).

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 If you are creating the object for an SSL VPN hosted on an IOS device, you can enter a name for the heading that is displayed above the bookmarks list in the **Bookmarks Heading (IOS)** field.

Step 6 The Bookmarks table displays any URLs that are defined for the object. To add a bookmark, click the **Add Row** button below the table; to edit an existing bookmark, select it and click the **Edit Row** button.

The Add/Edit SSL VPN Bookmark Entry dialog box opens. For more information about the fields on this dialog box, see Add or Edit Bookmark Entry Dialog Boxes, page 34-45.

- In the **Bookmark Option** field, select whether you are defining a bookmark (**Enter Bookmark**) or adding bookmarks from another SSL VPN bookmark object (**Include Existing Bookmarks**). If you are including an existing object, enter the object's name or click **Select** to select it from a list of existing objects.

- If you are creating the object for use on an IOS device, enter the title of the bookmark, which is displayed to users, and the URL. Be careful to select the correct protocol for the URL. Click **OK** to add the bookmark to the table of bookmarks.
- If you are creating the object for use on an ASA device, you have many more options. Besides the title and the URL, you can define a subtitle and image icon for the bookmark plus other options.

Tip If you choose the protocols RDP, SSH, Telnet, VNC, or ICA, you must configure the plug-in for the protocol in the **Remote Access VPN > SSL VPN > Other Settings** policy (see Configuring SSL VPN Browser Plug-ins (ASA), page 31-60).

You can also configure the bookmark to use the Post method rather than the Get method. If you use Post, you must configure the post parameters by clicking **Add Row** beneath the Post Parameters table. For more information on Post parameters, see these topics:

- Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84
- Add and Edit Post Parameter Dialog Boxes, page 34-48

Click **OK** to add the bookmark to the table of bookmarks.

Step 7 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 8 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 9 Click **OK** to save the object.

Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks

One of the options you have for configuring bookmarks on an SSL VPN hosted on an ASA device is the method used by a URL, either Get or Post. The Get method is the standard method; a user clicks the URL and is taken to the web page. The Post method is useful when processing the data might involve changes to it, for example, storing or updating data, ordering a product, or sending e-mail.

If you choose the Post URL method, you must configure Post parameters for bookmark entries. Because these are often personalized resources that contain the user ID and password or other input parameters, you might need to define clientless SSL VPN macro substitutions.

Clientless SSL VPN macro substitutions let you configure users for access to personalized resources that contain the user ID and password or other input parameters. Examples of such resources include bookmark entries, URL lists, and file shares.

Note For security reasons, password substitutions are disabled for file access URLs (cifs://). Also for security reasons, use caution when introducing password substitutions for web links, especially for non-SSL instances.

You can use the following macro substitutions:

- **Logon Information Substitutions**— The security appliance obtains values for these substitutions from the SSL VPN Logon page. It recognizes these strings in user requests, and replaces them with the value specific to the user before it passes the request on to a remote server.

These are the available macro substitutions:

- CSCO_WEBVPN_USERNAME

  The username used to log into the SSL VPN.

- CSCO_WEBVPN_PASSWORD

  The password used to log into the SSL VPN.

- CSCO_WEBVPN_INTERNAL_PASSWORD

  The internal resource password entered when logging into the SSL VPN.

- CSCO_WEBVPN_CONNECTION_PROFILE

  The connection profile associated with the user group selected when logging into the SSL VPN.

For example, if a URL list contains the link http://someserver/homepage/CSCO_WEBVPN_USERNAME.html, the security appliance translates it to the following unique links:

- For USER1 the link becomes http://someserver/homepage/USER1.html
- For USER2 the link is http://someserver/homepage/USER2.html

In the following example, cifs://server/users/CSCO_WEBVPN_USERNAME lets the security appliance map a file drive to specific users:

- For USER1 the link becomes cifs://server/users/USER1
- For USER2 the link is cifs://server/users/USER2

- **RADIUS/LDAP Vendor-Specific Attributes (VSAs)**—These substitutions let you set substitutions configured on either a RADIUS or an LDAP server. These are the available macro substitutions:
  - CSCO_WEBVPN_MACRO1
  - CSCO_WEBVPN_MACRO2

For information on configuring bookmarks, see Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82.

Configuring SSL VPN Smart Tunnels for ASA Devices

A smart tunnel is a connection between an application running on a user's workstation and a private site. The connection uses a clientless (browser-based) SSL VPN session with the security appliance as the pathway and proxy server. Smart tunnels do not require the user to connect the application to the local port, so the application can gain access to the network without giving the user administrative privileges, as is required for full tunnel support. If you do not otherwise configure the network to allow access to an application, you can create a smart tunnel for those applications that you want to support.

You can configure smart tunnel access to an application under the following conditions:

- The application is a Winsock 2, TCP-based application and there is a browser plug-in for the application. Cisco distributes plug-ins for some applications for use in clientless SSL VPN, including SSH (for both SSH and Telnet sessions), RDP, and VNC. You must supply or obtain plug-ins for any other applications. Configure plug-ins in the **Remote Access VPN > SSL VPN > Other Settings** policy on the Plug-Ins tab.
- The user's workstation is a supported platform. See the Cisco ASA 5500 Series Adaptive Security Appliances documentation that corresponds with your ASA version for supported platforms, http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html.

Users of Microsoft Windows Vista who use smart tunnels (or port forwarding) must add the URL of the ASA device to the Trusted Site zone. Configure the Trusted Site zone in Internet Explorer (**Tools > Internet Options**, **Security** tab).

- The user's browser must be enabled with Java, Microsoft ActiveX, or both.
- If the user's workstation requires a proxy server to reach the security appliance, the URL of the terminating end of the connection must be in the list of URLs excluded from proxy services. In this configuration, smart tunnels support only basic authentication.

Tip A stateful failover does not retain smart tunnel connections. Users must reconnect following a failover.

You configure smart tunnel access for an application by creating an SSL VPN smart tunnel list policy object and including that object in an ASA group policy object. You then assign the ASA group policy object to a device in the **Remote Access VPN > Group Policies** policy.

Related Topics

- Understanding Group Policies (ASA), page 31-27
- Creating Policy Objects, page 6-9
- Policy Object Manager, page 6-4

Step 1 Create an SSL VPN smart tunnel list policy object:

a. Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4), and select **SSL VPN Smart Tunnel Lists** from the table of contents.

Tip You can also create SSL VPN smart tunnel list objects when you create or edit the ASA group policy object. For more information, see Selecting Objects for Policies, page 6-2.

b. Click the **Add Object** button to open the Add and Edit Smart Tunnel List Dialog Boxes, page 34-66.

c. Enter a name for the object, up to 64 characters.

d. To the table of applications, add those applications for which you are granting smart tunnel access (click the **Add Row** button to open the Add and Edit A Smart Tunnel Entry Dialog Boxes, page 34-67). Consider the following:

- Enter an application name that is easy to understand and include version numbers if you support more than one version. For example, Microsoft Outlook.
- For the application path, entering only the filename, for example, outlook.exe, is the simplest and most maintainable option. This allows the user to install the application in any folder. Enter the full path if you want to enforce a specific installation structure.
- Hash values are optional, but you can use them to prevent spoofing. Without hash values, a user can rename an application to a supported filename; the security appliance checks only the filename and path (if specified). However, if you enter hash values, you must maintain them as users apply patches or application upgrades. For specific information on determining hash values, see Add and Edit A Smart Tunnel Entry Dialog Boxes, page 34-67.

Click **OK** to save the entry.

e. You can also incorporate other SSL VPN smart list objects into the object. This allows you to create a core set of smart list objects that you can use repeatedly in other objects.

f. Click **OK** to save the object.

Step 2 (Optional) Create an SSL VPN smart tunnel auto sign-on list policy object:

a. Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4), and select **SSL VPN Smart Tunnel Auto Signon Lists** from the table of contents.

Tip You can also create SSL VPN smart tunnel auto sign-on list objects when you create or edit the ASA group policy object. For more information, see Selecting Objects for Policies, page 6-2.

b. Click the **Add Object** button to open the Add and Edit Smart Tunnel Auto Signon List Dialog Boxes, page 34-71.

c. Enter a name for the object, up to 64 characters.

d. To the table of smart tunnel auto sign-on entries, add the servers for which to automate the submission of login credentials during smart tunnel setup (click the **Add Row** button to open the Add and Edit Smart Tunnel Auto Signon Entry Dialog Boxes, page 34-72).

e. You can also incorporate other SSL VPN smart tunnel auto sign-on list objects into the object. This allows you to create a core set of smart tunnel auto sign-on list objects that you can use repeatedly in other objects.

f. Click **OK** to save the object.

Step 3 Configure the ASA group policy object to use the SSL VPN smart tunnel list object:

a. Edit (or create) the ASA group policy object either from the Policy Object Manager, page 6-4 or the **Remote Access VPN > Group Policies** policy. The object must be configured to support SSL VPNs. (You can also edit these objects from the **Remote Access VPN > Connection Profiles** policy from an individual profile.)

b. Select the **SSL VPN > Clientless** folder from the table of contents to open ASA Group Policies SSL VPN Clientless Settings, page 34-12.

c. Enter the name of the SSL VPN smart tunnel list object in the **Smart Tunnel** field.

d. Select **Auto Start Smart Tunnel** to automatically start smart tunnels for the applications when the user connects to the SSL VPN portal.

If you do not select this option, users must start smart tunnel access using the **Application Access > Start Smart Tunnels** button on the clientless SSL VPN portal page.

e. Enter the name of the SSL VPN smart tunnel auto sign-on list object in the **Smart Tunnel Auto Signon Server List** field.

f. If the universal naming convention (domain\username) is required for authentication, specify the Windows domain to add it to the username during auto sign-on in the **Domain Name** field. For example, enter CISCO to specify CISCO\qa_team when authenticating for the username qa_team. You must also check the Use Domain option when configuring associated entries in the auto sign-on server list.

Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Clientless SSL VPN uses WINS and the Common Internet File System (CIFS) protocol to access or share files, printers, and other machine resources on remote systems. The ASA or IOS device uses a proxy CIFS client to provide this access transparently; users appear to have direct access to the file systems (subject to individual file and user permissions).

When users attempt a file-sharing connection to a Windows computer by using its computer name, the file server they specify corresponds to a specific WINS name that identifies a resource on the network. The security appliance queries WINS or NetBIOS name servers to map WINS names to IP addresses. SSL VPN requires NetBIOS to access or share files on remote systems.

You use WINS server list policy objects to configure the list of WINS servers that are used to resolve these Microsoft file-directory share names. The WINS server list objects define the NetBIOS Name Service (NBNS) server list on the device (using the **nbns-list** and **nbns-server** commands) for Common Internet File System (CIFS) name resolution.

After creating the WINS server list policy object, you can configure it in the following policies and policy objects, and also select the file access services that you want to allow:

- ASA devices—In the **Remote Access VPN > Connection Profiles** policy, specify the WINS server list object on the **SSL** tab (see SSL Tab (Connection Profiles), page 31-22).

  Select the file access options on the **SSL VPN > Clientless** page in an ASA group policy object (see ASA Group Policies SSL VPN Clientless Settings, page 34-12), which you then select in one of these policies:

  - **Remote Access VPN > Group Policies**
  - **Remote Access VPN > Connection Profiles** on the **General** tab

- IOS devices—On the **Clientless** page in a user group policy object configured for SSL VPN (see User Group Dialog Box—Clientless Settings, page 34-83), which you then select in the **Remote Access VPN > SSL VPN** policy on the **General** tab.

Related Topics

- Creating Policy Objects, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager, page 6-4.

Tip You can also create WINS server list objects when defining policies or objects that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

Step 2 Select **WINS Server Lists** from the Object Type selector.

The WINS Server List page opens, displaying the currently defined WINS server list objects.

Step 3 Right-click in the work area and select **New Object** to open the Add or Edit WINS Server List Dialog Box, page 34-89.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Click the **Add Row** button below the table, or select a server in the table and click **Edit Row**, to configure the WINS servers defined in the object. Configure these settings:

- **Server**—The IP address of the WINS server. You can select a network/host object or enter the address directly.
- **Set as Master Browser**—Select this option if the server is a primary browser, which maintains the list of computers and shared resources.

Other fields are optional; change them if you want non-default values. For more information, see Add or Edit WINS Server Dialog Box, page 34-90.

Click **OK** to save your changes.

Step 6 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 7 (Optional) Select **Allow Value Override per Device** to allow the properties of this object to be redefined on individual devices. See Allowing a Policy Object to Be Overridden, page 6-18.

Step 8 Click **OK** to save the object.

Managing Dynamic Access Policies for Remote Access VPNs (ASA 8.0+ Devices)

This chapter explains Dynamic Access Policies (DAP) for assigning remote access users to connection profiles (tunnel groups). You can configure these policies for remote access IKEv1 IPsec on ASA 8.0+ devices, IKEv2 IPsec on ASA 8.4(x) devices, and SSL VPNs on ASA 8.0+ (except 8.5) devices.

For information on configuring other remote access policies for ASA and PIX 7.0+ devices, see Chapter 31, "Managing Remote Access VPNs on ASA and PIX 7.0+ Devices".

This chapter contains the following topics:

- Understanding Dynamic Access Policies, page 32-1
- Configuring Dynamic Access Policies, page 32-2
- Dynamic Access Page (ASA), page 32-11

Understanding Dynamic Access Policies

Multiple variables can affect each VPN connection, for example, intranet configurations that frequently change, the various roles each user may inhabit within an organization, and logins from remote access sites with different configurations and levels of security. The task of authorizing users is much more complicated in a VPN environment than it is in a network with a static configuration.

Dynamic access policies (DAP) on a security appliance let you configure authorization that addresses these many variables. You create a dynamic access policy by setting a collection of access control attributes that you associate with a specific user tunnel or session. These attributes address issues of multiple group membership and endpoint security. That is, the security appliance grants access to a particular user for a particular session based on the policies you define. It generates a DAP at the time the user connects by selecting and/or aggregating attributes from one or more DAP records. It selects these DAP records based on the endpoint security information of the remote device and the AAA authorization information for the authenticated user. It then applies the DAP record to the user tunnel or session. The DAP system includes the following components that require your attention:

- DAP Selection Configuration File—A text file containing criteria that the security appliance uses for selecting and applying DAP records during session establishment. It is stored on the security appliance. You can use Security Manager to modify it and upload it to the security appliance in XML data format. DAP selection configuration files include all of the attributes that you configure. These can include AAA attributes, endpoint attributes, and access policies as configured in network and web-type ACL filter, port forwarding, and URL lists.

- DfltAccess Policy—Always the last entry in the DAP summary table, always with a priority of 0. You can configure Access Policy attributes for the default access policy, but it does not contain—and you cannot configure—AAA or endpoint attributes. You cannot delete the DfltAccessPolicy, and it must be the last entry in the summary table.

Tip Dynamic Access policies take precedence over Group policies. If a setting is not specified in a Dynamic Access policy, an ASA device checks for Group policies that specify the setting.

Integration of Cisco Secure Desktop with DAP

The security appliance integrates the Cisco Secure Desktop (CSD) features into dynamic access policies (DAPs). Depending on the configuration, the security appliance uses one or more endpoint attribute values in combination with optional, AAA attribute values as conditions for assigning a DAP. The Cisco Secure Desktop features supported by the endpoint attributes of DAPs include OS detection, prelogin policies, Basic Host Scan results, and Endpoint Assessment.

As an administrator, you can specify a single attribute or combine attributes that together form the conditions required to assign a DAP to a session. The DAP provides network access at the level that is appropriate for the endpoint AAA attribute value. The security appliance applies a DAP when all of its configured endpoint criteria are satisfied.

Related Topics

- Configuring Dynamic Access Policies, page 32-2
- Configuring DAP Attributes, page 32-7

Configuring Dynamic Access Policies

This procedure describes how to create or edit a dynamic access policy.

Related Topics

- Understanding Dynamic Access Policies, page 32-1
- Understanding DAP Attributes, page 32-4
- Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > Dynamic Access** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Dynamic Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Dynamic Access page opens. For a description of the elements on this page, see Dynamic Access Page (ASA), page 32-11.

Step 2 Click **Create** or select a policy in the table and click **Edit**.

The Add/Edit Dynamic Access Policy dialog box opens, with the Main tab open by default. For a description of the elements in this dialog box, see Table 32-4 on page 32-13.

Step 3 Enter the name of the DAP record (up to 128 characters).

Step 4 Specify a priority for the DAP record. The security appliance applies access policies in the order you set here, highest number having the highest priority.

Step 5 Enter a description for the DAP record.

Step 6 In the **Main** tab, configure the DAP attributes and the type of remote access method supported by the DAP system on your security appliance. For a detailed description of the elements on this tab, see Table 32-5 on page 32-14.

a. Click **Create** below the table, or select a DAP entry in the table and click **Edit**. The Add/Edit DAP Entry dialog box opens. For a description of the elements on this dialog box, see Add/Edit DAP Entry Dialog Box, page 32-21.

   For a full description of the procedure to define the DAP attributes, see Configuring DAP Attributes, page 32-7.

b. Select the type of remote access permitted by the DAP system.

c. Select the **Network ACL** tab to select and configure network ACLs to apply to this DAP record. Beginning with Security Manager version 4.10, you can select Unified ACL entries in addition to Extended entries.

   This tab is available only if you selected an access method other than Web Portal.

d. Select the **WebType ACL** tab to select and configure Web-type ACLs to apply to this DAP record.

   This tab is available only if you selected an access method other than AnyConnect Client.

e. Select the **Functions** tab to configure file server entry and browsing, HTTP proxy, and URL entry for the DAP record.

   This tab is available only if you selected an access method other than AnyConnect Client.

f. Select the **Port Forwarding** tab to select and configure port forwarding lists for user sessions.

   This tab is available only if you selected an access method other than AnyConnect Client.

g. Select the **Bookmark** tab to select and configure URL lists for user sessions.

   This tab is available only if you selected an access method other than AnyConnect Client.

h. Select the **Action** tab to configure the type of remote access permitted.

   This tab is available for all types of access methods.

i. Select the **AnyConnect** tab to choose if the setting for Always-On VPN on the AnyConnect service profile remains unchanged, is disabled, or the AnyConnect Profile Setting must be used. Always-On VPN enables AnyConnnect to automatically establish a VPN session after you log onto the system.

j. Select the **Custom Attributes** tab to add AnyConnect Custom Attributes.

   This tab is available only if you selected the access method as Unchanged, AnyConnect Client, Both Default Web Portal, or Both Default Anyconnect Client. For information about how to add AnyConnect Custom Attributes, see Add/Edit AnyConnect Custom Attribute Dialog Box, page 31-71.

Step 7 Select the **Logical Operations** tab to create multiple instances of each type of endpoint attribute. For a description of the elements on this tab, see Table 32-22 on page 32-43.

Step 8 Select the **Advanced Expressions** tab to set additional attributes for the DAP using free-form LUA. For a description of the elements on this tab, see Table 32-23 on page 32-45.

Step 9 Click **OK**.

Understanding DAP Attributes

DAP records include all of the attributes that you configure. These can include AAA attributes, endpoint attributes, and access policies as configured in network and web-type ACL filter, port forwarding and URL lists.

DAP and AAA Attributes

DAP complements AAA services. It provides a limited set of authorization attributes that can override those AAA provides. The security appliance selects DAP records based on the AAA authorization information for the user and posture assessment information for the session. The security appliance can select multiple DAP records depending on this information, which it then aggregates to create DAP authorization attributes.

You can specify AAA attributes from the Cisco AAA attribute hierarchy, or from the full set of response attributes that the security appliance receives from a RADIUS or LDAP server.

AAA Attribute Definitions

Table 32-1 on page 32-4 defines the AAA selection attribute names that are available for DAP use. The Attribute Name field shows you how to enter each attribute name in a LUA logical expression, which you might do on the Advanced tab of the Add/Edit Dynamic Access Policy dialog box.

Table 32-1 AAA Attribute Definitions

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Cisco	aaa.cisco.memberof	AAA	string	128	memberof value
	aaa.cisco.username	AAA	string	64	username value
	aaa.cisco.class	AAA	string	64	class attribute value
	aaa.cisco.ipaddress	AAA	number	–	framed-ip address value
	aaa.cisco.tunnelgroup	AAA	string	64	tunnel-group name
LDAP	aaa.ldap.<label>	LDAP	string	128	LDAP attribute value pair
RADIUS	aaa.radius.<number>	RADIUS	string	128	Radius attribute value pair

DAP and Endpoint Security

The security appliance obtains endpoint security attributes by using posture assessment methods that you configure. These include Cisco Secure Desktop and NAC. You can use a match of a prelogin policy, Basic Host Scan entry, Host Scan Extension, or any combination of these and any other policy attributes to assign access rights and restrictions. At minimum, configure DAPs to assign to each prelogin policy and Basic Host Scan entry.

Endpoint Assessment, a Host Scan extension, examines the remote computer for a large collection of antivirus and antispyware applications, associated definitions updates, and firewalls. You can use this feature to combine endpoint criteria to satisfy your requirements before the security appliance assigns a specific DAP to the session.

DAP and Anti-Virus, Anti-Spyware, and Personal Firewall Programs

The security appliance uses a DAP policy when the user attributes matches the configured AAA and endpoint attributes. The Prelogin Assessment and Host Scan modules of Cisco Secure Desktop return information to the security appliance about the configured endpoint attributes, and the DAP subsystem uses that information to select a DAP record that matches the values of those attributes. Most, but not all, anti-virus, anti-spyware, and personal firewall programs support active scan, which means that the programs are memory-resident, and therefore always running. Host Scan checks to see if an endpoint has a program installed, and if it is memory-resident as follows:

- If the installed program does not support active scan, Host Scan reports the presence of the software. The DAP system selects DAP records that specify the program.
- If the installed program does support active scan, and active scan is enabled for the program, Host Scan reports the presence of the software. Again the security appliance selects DAP records that specify the program.
- If the installed program does support active scan and active scan is disabled for the program, Host Scan ignores the presence of the software. The security appliance does not select DAP records that specify the program. Further, the output of the **debug trace** command, which includes a lot of information about DAP, does not indicate the program presence, even though it is installed.

Endpoint Attribute Definitions

Table 32-2 on page 32-5 defines the endpoint selection attribute names that are available for DAP use. The Attribute Name field shows you how to enter each attribute name in a LUA logical expression, which you might do on the Advanced tab of the Add/Edit Dynamic Access Policy dialog box. The *label* variable identifies the application, filename, process, or registry entry.

Table 32-2 Endpoint Attribute Definitions

Attribute Type	Attribute Name	Source	Value	Max String Length	Description
Antispyware (Requires Cisco Secure Desktop)	endpoint.as.label.exists	Host Scan	true	–	Antispyware program exists
	endpoint.as.label.version		string	32	Antispyware description
	endpoint.as.label.description		string	128	class attribute value
	endpoint.as.label.lastupdate		integer	–	Seconds since update of antispyware definitions
Antivirus (Requires Cisco Secure Desktop)	endpoint.av.label.exists	Host Scan	true	–	Antivirus program exists
	endpoint.av.label.version		string	32	Antivirus description
	endpoint.av.label.description		string	128	class attribute value
	endpoint.av.label.lastupdate		integer	–	Seconds since update of antivirus definitions

Table 32-2 Endpoint Attribute Definitions (continued)

Application	endpoint.application.clienttype	Application	string	–	Client type: CLIENTLESS ANYCONNECT IPSEC L2TP
File	endpoint.file.label.exists	Secure Desktop	true	–	The files exists
	endpoint.file.label.last modified		integer	–	Seconds since file was last modified
	endpoint.file.label.crc.32		integer	–	CRC32 hash of the file
NAC	endpoint.nac.status	NAC	string	-	User defined status string
Operating System	endpoint.os.version	Secure Desktop	string	32	Service pack for Windows
	endpoint.os.servicepack		integer	–	Operating system
Personal firewall (Requires Secure Desktop)	endpoint.fw.label.exists	Host Scan	true	–	The personal firewall exists
	endpoint.fw.label.version		string	32	Version
	endpoint.fw.label.description		string	128	Personal firewall description
Policy	endpoint.policy.location	Secure Desktop	string	64	Location value from Cisco Secure Desktop
Process	endpoint.process.label.exists	Secure Desktop	true	–	The process exists
	endpoint.process.label.path		string	255	Full path of the process
Registry	endpoint.registry.label.type	Secure Desktop	dword string	–	dword
	endpoint.registry.label.value		string	255	Value of the registry entry
VLAN	endpoint.vlan.type	CNA	sting	–	VLAN type: ACCESS AUTH ERROR GUEST QUARANTINE ERROR STATIC TIMEOUT

About Advanced Expressions for AAA or Endpoint Attributes

In the text box you enter free-form LUA text that represents AAA and/or endpoint selection logical operations. ASDM does not validate text that you enter here; it just copies this text to the DAP policy file, and the security appliance processes it, discarding any expressions it cannot parse.

This option is useful for adding selection criteria other than what is possible in the AAA and endpoint attribute areas above. For example, while you can configure the security appliance to use AAA attributes that satisfy any, all, or none of the specified criteria, endpoint attributes are cumulative, and must all be satisfied. To let the security appliance employ one endpoint attribute or another, you need to create appropriate logical expressions in LUA and enter them here.

Examples of DAP Logical Expressions

Study these examples for help in creating logical expressions in LUA.

- This AAA LUA expression tests for a match on usernames that begin with "b". It uses the string library and a regular expression:

```
not(string.find(aaa.cisco.username, "^b") == nil)
```

- This endpoint expression tests for a match on CLIENTLESS OR CVC client types:

```
endpoint.application.clienttype=="CLIENTLESS" or endpoint.application.clienttype=="CVC"
```

- This endpoint expression tests for Norton Antivirus versions 10.x but excludes 10.5.x:

```
(endpoint.av.NortonAV.version > "10" and endpoint.av.NortonAV.version < "10.5") or
endpoint.av.NortonAV.version > "10.6"
```

DAP Connection Sequence

The following sequence outlines a typical remote access connection establishment.

1. A remote client attempts a VPN connection.
2. The security appliance performs posture assessment, using configured NAC and Cisco Secure Desktop Host Scan values.
3. The security appliance authenticates the user via AAA. The AAA server also returns authorization attributes for the user.
4. The security appliance applies AAA authorization attributes to the session, and establishes the VPN tunnel.
5. The security appliance selects DAP records based on the user AAA authorization information and the session posture assessment information.
6. The security appliance aggregates DAP attributes from the selected DAP records, and they become the DAP policy.
7. The security appliance applies the DAP policy to the session.

Related Topics

- Configuring Dynamic Access Policies, page 32-2
- Understanding Dynamic Access Policies, page 32-1
- Configuring DAP Attributes, page 32-7

Configuring DAP Attributes

The attributes you must define for a DAP policy include specifying the authorization attributes and endpoint attributes. You can also configure network and webtype ACLs, file browsing, file server entry, HTTP proxy, URL entry, port forwarding lists and URL lists.

This procedure describes how to create or edit the AAA and endpoint attributes required for a DAP policy.

Related Topics

- Understanding DAP Attributes, page 32-4
- Understanding Dynamic Access Policies, page 32-1
- Configuring Dynamic Access Policies, page 32-2

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > Dynamic Access** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Dynamic Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Dynamic Access page opens. For a description of the elements on this page, see Dynamic Access Page (ASA), page 32-11.

Step 2 Click **Create** on the Dynamic Access policy page, or select the row of a policy in the table on the page, and click **Edit**.

The Add/Edit Dynamic Access Policy dialog box opens, displaying the Main tab. For a description of the elements on the Main tab, see Main Tab, page 32-14.

Step 3 Click **Create** below the table, or select a DAP entry in the table and click **Edit**. The Add/Edit DAP Entry dialog box opens. For a description of the elements on this dialog box, see Add/Edit DAP Entry Dialog Box, page 32-21.

Step 4 Select the attribute type from the Criterion list, then enter the appropriate values. The dialog box values vary based on your selection. Options are:

- AAA Attributes Cisco; see Table 32-6 on page 32-23.
- AAA Attributes LDAP; see Table 32-7 on page 32-25.
- AAA Attributes RADIUS; see Table 32-8 on page 32-26.
- Anti-Spyware; see Table 32-9 on page 32-27.
- Anti-Virus; see Table 32-10 on page 32-28.
- AnyConnect Identity; see Table 32-11 on page 32-29.
- Application; see Table 32-12 on page 32-30.
- Device; see Table 32-13 on page 32-31.
- File; see Table 32-14 on page 32-32.
- NAC; see Table 32-15 on page 32-33.
- Operating System; see Table 32-16 on page 32-34.
- Personal Firewall; see Table 32-17 on page 32-35.
- Policy; see Table 32-18 on page 32-36.
- Process; see Table 32-19 on page 32-36.
- Registry; see Table 32-20 on page 32-38.
- Multiple Certificate Authentication; see Table 32-21 on page 32-39

Step 5 Click **OK**.

Configuring Cisco Secure Desktop Policies on ASA Devices

Cisco Secure Desktop (CSD) provides a reliable means of eliminating all traces of sensitive data by providing a single, secure location for session activity and removal on the client system. CSD provides a session-based interface where sensitive data is shared only for the duration of an SSL VPN session. All session information is encrypted, and all traces of the session data are removed from the remote client when the session is terminated, even if the connection terminates abruptly. This ensures that cookies, browser history, temporary files, and downloaded content do not remain on a system.

When the session closes, CSD overwrites and removes all data using a U.S. Department of Defense (DoD) sanitation algorithm to provide endpoint security protection.

Note A complete explanation of the capabilities and configuration of the Cisco Secure Desktop program is beyond the scope of this document. For information about configuring CSD, and what CSD can do for you, see the materials available online at http://www.cisco.com/en/US/products/ps6742/tsd_products_support_configure.html. Select the configuration guide for the CSD version you are configuring.

This procedure describes how to configure the Cisco Secure Desktop feature on an ASA device.

Before You Begin

- Make sure a connection profile policy has been configured on the device. See Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5

Step 1 Do one of the following:

- (Device view) With an ASA device selected, select **Remote Access VPN > Dynamic Access** from the Policy selector.
- (Policy view) Select **Remote Access VPN > Dynamic Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

The Dynamic Access page opens. For a description of the elements on this page, see Dynamic Access Page (ASA), page 32-11.

Step 2 In the Cisco Secure Desktop section, select **Enable CSD** to enable CSD on the ASA device.

Note The Enable CSD option is available for devices running ASA version less than ASA 9.5(2). Beginning with Security Manager 4.10, a new check box is available to configure Hostscan (to disable CSD) only for devices running the ASA version 9.5(2) or later.

Step 3 In the **CSD Package** field, specify the name of the File Object that identifies the Cisco Secure Desktop package you want to upload to the device. Click **Select** to select an existing File Object or to create a new one. For more information, see Add and Edit File Object Dialog Boxes, page 34-37.

Note The package version must be compatible with the ASA operating system version. When you create a local policy in Device view, the **Version** field indicates the CSD package version you should select. (The version is included in the package file name. For example, securedesktop-asa_k9-3.3.0.118.pkg is CSD version 3.3.0.118.) When you create a shared policy in Policy view, the **Version** field indicates the version of the CSD file you selected. For more information on version compatibility, see Understanding and Managing SSL VPN Support Files, page 30-5.

Step 4 (Optional) In the **Hostscan Package** field, specify the name of the File Object that identifies the Host Scan package you want to upload to the device. Click **Select** to select an existing File Object or to create a new one. For more information, see Add and Edit File Object Dialog Boxes, page 34-37.

Step 5 Click **Configure** to open the Cisco Secure Desktop Manager (CSDM) Policy Editor that lets you configure CSD on the security appliance. This application is independent of Security Manager; read the CSD documentation cited above for an explanation of how to use the policy editor.

The editor contains these main items (select them in the table of contents):

- Prelogin Policies—This is a decision tree. When a user attempts a connection, the user's system is evaluated against your rules and the first rule that matches is applied. Typically, you create policies for secure locations, home locations, and insecure public locations. You can make your checks based on registry information, the presence of specific files or certificates, the workstation's operating system, or IP address.

  All editing is done through the right-click menu. Right click on boxes or + signs to activate related settings, if any.

  For end nodes, you can select these options:

  - Access Denied—Workstations that match your criteria are prevented from accessing the network.
  - Policy—You want to define a specific admission policy at this point. After naming the policy, it is added to the table of contents. Select each item in the policy and configure its settings.
  - Subsequence—You want to perform additional checks. Enter the name of the next decision tree that you want to evaluate for this workstation.
- Host scan—You can specify a set of registry entries, file names, and process names, which form a part of the basic host scan. The host scan occurs after the prelogin assessment but before the assignment of a dynamic access policy. Following the basic host scan, the security appliance uses the login credentials, the host scan results, prelogin policy, and other criteria you configure to assign a dynamic access policy. You can also enable:
  - Endpoint Assessment—The remote workstation scans for a large collection of antivirus, antispyware, and personal firewall applications, and associated updates.
  - Advanced Endpoint Assessment—Includes all of the Endpoint Assessment features, and lets you configure an attempt to update noncompliant workstations to meet the version requirements you specify. You must purchase and install a license for this feature before you can configure it.

Dynamic Access Page (ASA)

Use the Dynamic Access page to view the dynamic access policies (DAP) defined on the security appliance. From this page, you can create, edit, or delete DAPs.

Use the Cisco Secure Desktop section to enable and download the Cisco Secure Desktop (CSD) software on the selected ASA device. Cisco Secure Desktop provides a single, secure location for session activity and removal on the client system, ensuring that sensitive data is shared only for the duration of an SSL VPN session.

Note The CSD client software must be installed and activated on a device in order for an SSL VPN policy to work properly.

Tip Dynamic Access policies take precedence over Group policies. If a setting is not specified in a Dynamic Access policy, an ASA device checks for Group policies that specify the setting.

Navigation Path

- (Device View) Select an ASA device; then select **Remote Access VPN > Dynamic Access (ASA)** from the Policy selector.
- (Policy View) Select **Remote Access VPN > Dynamic Access (ASA)** from the Policy Type selector. Select an existing policy or create a new one.

Related Topics

- Understanding Dynamic Access Policies, page 32-1
- Configuring Dynamic Access Policies, page 32-2
- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9

Field Reference

Table 32-3 Dynamic Access Policy Page (ASA)

Element	Description
Priority	Priority of the configured dynamic access policy record.
Name	Name of the configured dynamic access policy record.
Network ACL	Name of the firewall ACL that applies to the session.
WebType ACL	Name of the WebType VPN ACL that applies to the session.
Port Forwarding	Name of the port forwarding list that applies to the session.
Bookmark	Name of the SSL VPN Bookmark object that applies to the session.
Terminate	Indicates whether the session is terminated or not.
Description	Additional information about the configured dynamic access policy.
Create button	Click this button to create a dynamic access policy. See Add/Edit Dynamic Access Policy Dialog Box, page 32-12.

Table 32-3 Dynamic Access Policy Page (ASA) (continued)

Element	Description
Edit button	Click this button to edit the selected dynamic access policy. See Add/Edit Dynamic Access Policy Dialog Box, page 32-12.
Delete button	Click this button to delete the selected dynamic access policies.
Cisco Secure Desktop For the procedure to configure CSD on an ASA device, see Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9.	
Enable CSD	When selected, enables the CSD on the device. Enabling CSD loads the specified Cisco Secure Desktop package. If you transfer or replace the CSD package file, disable and then enable CSD to load the file.
CSD Package	Specify the name of the File Object that identifies the Cisco Secure Desktop package you want to upload to the device. Click **Select** to select an existing File Object or to create a new one. For more information, see Add and Edit File Object Dialog Boxes, page 34-37.
Hostscan Package	Specify the name of the File Object that identifies the Hostscan package you want to upload to the device. Click **Select** to select an existing File Object or to create a new one. For more information, see Add and Edit File Object Dialog Boxes, page 34-37.
Version	The package version must be compatible with the ASA operating system version. When you create a local policy in Device view, the Version field indicates the CSD package version you should select. (The version is included in the package file name. For example, securedesktop-asa_k9-3.3.0.118.pkg is CSD version 3.3.0.118.) When you create a shared policy in Policy view, the Version field indicates the version of the CSD file you selected. For more information on version compatibility, see Understanding and Managing SSL VPN Support Files, page 30-5.
Configure	Click **Configure** to open the Cisco Secure Desktop Manager (CSDM) Policy Editor that lets you configure CSD on the security appliance. For a description of the elements in this dialog box, see Cisco Secure Desktop Manager Policy Editor Dialog Box, page 32-46.

Add/Edit Dynamic Access Policy Dialog Box

Use the Add/Edit Dynamic Access Policy dialog box to configure the dynamic access policies (DAP) on your security appliance. You can specify a name for the dynamic access policy that you are adding, select the priority, specify attributes in a LUA expression, and set attributes for network and webtype ACL filters, file access, HTTP proxy, URL entry and lists, port forwarding, and clientless SSL VPN access methods.

Note For detailed information about dynamic access policy attributes, see Understanding DAP Attributes, page 32-4

These tabs are available in the Add/Edit Dynamic Access Policy dialog box:

- Main Tab, page 32-14
- Logical Operations Tab, page 32-42
- Advanced Expressions Tab, page 32-44

Navigation Path

Open the Dynamic Access Page (ASA), page 32-11, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit Dynamic Access Policy dialog box is displayed.

Related Topics

- Understanding Dynamic Access Policies, page 32-1
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-4 Add/Edit Dynamic Access Policy Dialog Box

Element	Description
Name	The name of the dynamic access policy record (up to 128 characters).
Priority	A priority for the dynamic access policy record. The security appliance applies access policies in the order you set here, highest number having the highest priority. In the case of dynamic access policy records with the same priority setting and conflicting ACL rules, the most restrictive rule applies. Priority is supported by Security Manager version 4.12 onwards for Multi-Context ASA version 9.6(2) or later devices.
Description	Additional information about the dynamic access policy record (up to 1024 characters). Description is supported by Security Manager version 4.12 onwards for Multi-Context ASA version 9.6(2) or later devices.
Main tab	Enables you to add a dynamic access policy entry and set attributes for the access policy depending on the type of remote access that you configure. For a description of the elements on this tab, see Main Tab, page 32-14.
Logical Operations tab	Enables you to create multiple instances of each type of endpoint attribute. For a description of the elements on this tab, see Logical Operations Tab, page 32-42.
Advanced Expressions tab	Enables you to configure one or more logical expressions to set AAA or endpoint attributes other than what is possible in the AAA and Endpoint areas. For a description of the elements on this tab, see Advanced Expressions Tab, page 32-44.

Main Tab

Use the Main tab of the Add/Edit Dynamic Access Policy dialog box to configure the dynamic access policy attributes and the type of remote access method supported your security appliance. You can set attributes for network and webtype ACL filters, file access, HTTP proxy, URL entry and lists, port forwarding, and clientless SSL VPN access methods.

Navigation Path

The Main tab appears when you open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12.

Related Topics

- Configuring Dynamic Access Policies, page 32-2
- Configuring DAP Attributes, page 32-7

Field Reference

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab

Element	Description
Criteria ID	The AAA and endpoint selection attribute names that are available for dynamic access policy use.
Content	Values of the AAA and endpoint attributes criteria that the security appliance uses for selecting and applying a dynamic access policy record during session establishment. Attribute values that you configure here override authorization values in the AAA system, including those in existing group policy, tunnel group, and default group records.
Create button	Click this button to configure AAA and endpoint attributes as selection criteria for the DAP record. See Add/Edit DAP Entry Dialog Box, page 32-21.
Edit button	Click this button to edit the selected dynamic access policy. See Add/Edit DAP Entry Dialog Box, page 32-21.
Delete button	Click this button to delete the selected dynamic access policies.
Access Method	Specify the type of remote access permitted: • **Unchanged**—Continue with the current remote access method. • **AnyConnect Client**—Connect using the Cisco AnyConnect VPN Client. • **Web Portal**—Connect with clientless VPN. • **Both default Web Portal**—Connect via either clientless or the AnyConnect client, with a default of clientless. • **Both default AnyConnect Client**—Connect via either clientless or the AnyConnect client, with a default of AnyConnect.
Network ACL tab—Lets you select and configure network ACLs to apply to this dynamic access policy. An ACL for a dynamic access policy can contain permit or deny rules, but not both. If an ACL contains both permit and deny rules, the security appliance rejects it.	

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
Network ACL	Lists the Access Control Lists (ACLs) that will be used to restrict user access to the SSL†VPN. Beginning with Security Manager version 4.10, Network ACL supports IPv6 entries. Also IPv6 is supported for devices running the software version ASA 9.0 or later. This is applicable for both Network ACL and Web Type ACL. Click the **Select** button to open the Access Control Lists Selector from which you can make your selection. The ACL contains conditions that describe a traffic stream of packets, and actions that describe what should occur based on those conditions. Only ACLs having all permit or all deny rules are eligible. Network ACL is supported by Security Manager version 4.12 onwards for Multi-Context ASA version 9.6(2) or later devices.
AnyConnect tab—Lets you choose if the setting for Always-On VPN on the AnyConnect service profile remains unchanged, is disabled, or the AnyConnect Profile Setting must be used. Always-On VPN enables AnyConnnect to automatically establish a VPN session after you log onto the system.	
Custom Attributes tab—Lists the AnyConnect Custom Attribute Type and Custom Attribute Name. AnyConnect custom attributes allow for a more expeditious delivery and deployment of new endpoint features by giving the ASA the ability to generically support the addition of new client controls without the need for an ASA software upgrade. Beginning with version 4.7, Security Manager enables to add Custom Attribute Data to an existing Custom Attribute Type. This feature is supported for devices that are running the ASA software version 9.3(1) or higher.	
Attribute Type	Select the Attribute Type that you configured in Add/Edit AnyConnect Custom Attribute Dialog Box, page 31-71 page.
Attribute Name	Select the Attribute Name that you configured in Add/Edit AnyConnect Custom Attribute Data Dialog Box, page 31-71 page.
WebType ACL tab—Lets you select and configure web-type ACLs to apply to this dynamic access policy. An ACL for a dynamic access policy can contain only permit or deny rules. If an ACL contains both permit and deny rules, the security appliance rejects it.	
Web Type ACL	Specifies the WebType access control list that will be used to restrict user access to the SSL†VPN. Click the **Select** button to open the Access Control Lists Selector from which you can make your selection. Only ACLs having all permit or all deny rules are eligible. Beginning with version 4.10, you can enter IPv6 values for the Web Type ACL.
Functions tab—Lets you configure file server entry and browsing, HTTP proxy, and URL entry for the dynamic access policy.	

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
File Server Browsing	Specify the file server browsing setting to be configured on the portal page: • **Unchanged**—Uses values from the group policy that applies to this session. • **Enable**—Enables CIFS browsing for file servers or shared features. • **Disable**—Disables CIFS browsing for file servers or shared features. **Note** Browsing requires NBNS (Master Browser or WINS). If that fails or is not configured, we use DNS. The CIFS browse feature does not support internationalization.
File Server Entry	Specify the file server entry setting to be configured on the portal page: • **Unchanged**—Uses values from the group policy that applies to this session. • **Enable**—Enables a user from entering file server paths and names on the portal page. When enabled, places the file server entry drawer on the portal page. Users can enter pathnames to Windows files directly. They can download, edit, delete, rename, and move files. They can also add files and folders. Shares must also be configured for user access on the applicable Windows servers. Users might have to be authenticated before accessing files, depending on network requirements. • **Disable**—Disables a user from entering file server paths and names on the portal page.

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
HTTP Proxy	Specify how you want to configure the security appliance to terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers: • **Unchanged**—Uses values from the group policy that applies to this session. • **Enable**—Allows the forwarding of an HTTP applet proxy to the client. The proxy is useful for technologies that interfere with proper content transformation, such as Java, ActiveX, and Flash. It bypasses mangling while ensuring the continued use of the security appliance. The forwarded proxy modifies the browser's old proxy configuration and redirects all HTTP and HTTPS requests to the new proxy configuration. It supports virtually all client side technologies, including HTML, CSS, JavaScript, VBScript, ActiveX, and Java. The only browser it supports is Microsoft Internet Explorer. • **Disable**—Disables the forwarding of an HTTP applet proxy to the client. • **Auto-start**—Enables HTTP proxy and to have the DAP record automatically start the applets associated with these features.

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
URL Entry	Using SSL VPN does not ensure that communication with every site is secure. SSL VPN ensures the security of data transmission between the remote user's PC or workstation and the security appliance on the corporate network. If a user then accesses a non-HTTPS web resource (located on the Internet or on the internal network), the communication from the corporate security appliance to the destination web server is not secured. In a clientless VPN connection, the security appliance acts as a proxy between the end user web browser and target web servers. When a user connects to an SSL-enabled web server, the security appliance establishes a secure connection and validates the server SSL certificate. The end user browser never receives the presented certificate, so therefore cannot examine and validate the certificate. The current implementation of SSL VPN does not permit communication with sites that present expired certificates. Neither does the security appliance perform trusted CA certificate validation. Therefore, users cannot analyze the certificate an SSL-enabled web-server presents before communicating with it. Specify how the URL entry setting must be configured on the portal page: • **Unchanged**—Uses values from the group policy that applies to this session. • **Enable**—Allows a user from entering HTTP/HTTPS URLs on the portal page. If this feature is enabled, users can enter web addresses in the URL entry box, and use clientless SSL VPN to access those websites. • **Disable**—Disables a user from entering HTTP/HTTPS URLs on the portal page. **Note** To limit Internet access for users, select **Disable** for the URL Entry field. This prevents SSL VPN users from surfing the Web during a clientless VPN connection.

Port Forwarding tab—Lets you select and configure port forwarding lists for user sessions.

Note Port Forwarding does not work with some SSL/TLS versions.

Caution Make sure Sun Microsystems Java Runtime Environment (JRE) 1.4+ is installed on the remote computers to support port forwarding (application access) and digital certificates.

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
Port Forwarding	Select an option for the port forwarding lists that apply to this DAP record: • **Unchanged**—Removes the attributes from the running configuration. • **Enable**—Enables port forwarding on the device. • **Disable**—Disables port forwarding on the device. • **Auto-start**—Enables port forwarding, and to have the DAP record automatically start the port forwarding applets associated with its port forwarding lists.
Port Forwarding List	The Port Forwarding List, that defines the mapping of the port number on the client machine to the application's IP address and port behind the SSL VPN gateway. You can click **Select** to open the Port Forwarding List Selector from which you can select the required Port Forwarding List from a list of Port Forwarding List objects. A Port Forwarding List object defines the mappings of port numbers on the remote client to the application's IP address and port behind the SSL VPN gateway.
Bookmark tab—Lets you enable and configure SSL VPN bookmarks. When enabled, users who successfully log into the SSL VPN are presented with the portal page containing the list of defined bookmarks. These bookmarks enable users to access resources available on SSL VPN websites in Clientless access mode.	
Enable Bookmarks	Specify the file server browsing setting to be configured on the portal page: • **Unchanged**—Uses values from the group policy that applies to this session. • **Enable**—Enables bookmarks on the SSL VPN portal page. • **Disable**—Disables bookmarks on the SSL VPN portal page.
Bookmarks	A list of websites that will be displayed on the portal page as a bookmark to enable users to access the resources available on the SSL VPN websites. You can click **Select** to open the Bookmarks Selector from which you can select the required bookmark from a list or create a new bookmark, as desired.
Action tab—Specifies special processing to apply to a specific connection or session. Action Tab is supported by Security Manager version 4.12 onwards for Multi-Context ASA version 9.6(2) or later devices. Select one of the following options from the drop-down list:	
Continue	(Default) When selected, continues the session. By default, the access policy attributes are applied to the session and it is running.

Table 32-5 Add/Edit Dynamic Access Policy Dialog Box > Main Tab (continued)

Element	Description
Quarantine	When selected, quarantines the session. By selecting quarantine, you can restrict a particular client who already has an established tunnel through a VPN. Restricted ACLs are applied to a session to form a restricted group, based on the selected DAP record. When an endpoint is not compliant with an administratively defined policy, the user can still access services for remediation (such as updating the antivirus and so on), but restrictions are placed upon the user. After the remediation occurs, the user can reconnect, which invokes a new posture assessment. If this assessment passes, the user connects. **Note** This parameter requires an AnyConnect release that supports AnyConnect Secure Mobility features.
Terminate	When selected, terminates the session. By default, the access policy attributes are applied to the session and it is running.
User Message	Enter a text message to display on the portal page when this DAP record is selected. Maximum 128 characters. A user message displays as a yellow orb. When a user logs on it blinks three times to attract attention, and then it is still. If several DAP records are selected, and each of them has a user message, all user messages display. **Note** You can include in such messages URLs or other embedded text, which require that you use the correct HTML tags. For example: All contractors please read Instructions for the procedure to upgrade your antivirus software. **Note** User Message is supported from Security Manager version 4.12 for ASA devices running version 9.6(2) or later in Multi-context mode.

The supported Dynamic Access Policy CLIs in Security Manager version 4.12 onwards for Multi-Context ASA 9.6(2) devices are as follows:

- Dynamic-access-policy-record action
- description
- exit
- help
- network-acl
- no
- priority
- quit
- user-message

Add/Edit DAP Entry Dialog Box

Use the Add/Edit DAP Entry dialog box to specify the authorization attributes and endpoint attributes for a dynamic access policy. The security appliance selects the dynamic access policy based on the endpoint security information of the remote device and the AAA authorization information for the authenticated user. It then applies the dynamic access policy to the user tunnel or session.

For detailed information about dynamic access policy attributes, see Understanding DAP Attributes, page 32-4.

The content of the dialog box differs based on the criterion that you select. The criterion is the authorization or endpoint attribute that serves as the selection criterion that the security appliance uses for selecting and applying dynamic access policies during session establishment. You can select from the following criteria:

- AAA Attributes Cisco—Refers to user authorization attributes that are stored in the AAA hierarchical model. See Add/Edit DAP Entry Dialog Box > AAA Attributes Cisco, page 32-22
- AAA Attributes LDAP—Sets the LDAP client stores all native LDAP response attribute value pairs in a database associated with the AAA session for the user. See Add/Edit DAP Entry Dialog Box > AAA Attributes LDAP, page 32-24.
- AAA Attributes RADIUS—Sets the RADIUS client stores all native RADIUS response attribute value pairs in a database associated with the AAA session for the user. See Add/Edit DAP Entry Dialog Box > AAA Attributes RADIUS, page 32-25.
- Anti-Spyware—Creates an endpoint attribute of type Anti-Spyware. You can use the Host Scan modules of Cisco Secure Desktop to scan for antispyware applications and updates that are running on the remote computer. See Add/Edit DAP Entry Dialog Box > Anti-Spyware, page 32-26.
- Anti-Virus—Creates an endpoint attribute of type Anti-Virus. You can use the Host Scan modules of Cisco Secure Desktop to scan for antivirus applications and updates that are running on the remote computer. See Add/Edit DAP Entry Dialog Box > Anti-Virus, page 32-27.
- AnyConnect Identity—Creates an endpoint attribute of type AnyConnect Identity. See Add/Edit DAP Entry Dialog Box > AnyConnect Identity, page 32-28.
- Application—Indicates the type of remote access connection. See Add/Edit DAP Entry Dialog Box > Application, page 32-29.
- Device—Creates an endpoint attribute of type Device. The Device Criterion lets you provide specific device information for use during the associated prelogin policy checking. See Add/Edit DAP Entry Dialog Box > Device, page 32-30.
- File—Creates an endpoint attribute of type File. Filename checking to be performed by Basic Host Scan must be explicitly configured using Cisco Secure Desktop Manager. See Add/Edit DAP Entry Dialog Box > File, page 32-31.
- NAC—Creates an endpoint attribute of type NAC. NAC protects the enterprise network from intrusion and infection from worms, viruses, and rogue applications by performing endpoint compliancy. We refer to these checks as posture†validation. See Add/Edit DAP Entry Dialog Box > NAC, page 32-32.
- Operating System—Creates an endpoint attribute of type Operating System. The prelogin assessment module of the CSD can check the remote device for the OS version, IP address, and Microsoft Windows registry keys. See Add/Edit DAP Entry Dialog Box > Operating System, page 32-33.

- Personal Firewall—Creates an endpoint attribute of type Personal Firewall. You can use the Host Scan modules of Cisco Secure Desktop to scan for personal firewall applications and updates that are running on the remote computer. For a description of the elements in the dialog box, see Add/Edit DAP Entry Dialog Box > Personal Firewall, page 32-34.
- Policy—Creates an endpoint attribute of type Policy. See Add/Edit DAP Entry Dialog Box > Policy, page 32-35.
- Process—Process name checking to be performed by Basic Host Scan must be explicitly configured using Cisco Secure Desktop Manager. See Add/Edit DAP Entry Dialog Box > Process, page 32-36.
- Registry—Creates an endpoint attribute of type Registry. Registry key scans apply only to computers running Windows Microsoft Windows operating systems. See Add/Edit DAP Entry Dialog Box > Registry, page 32-37.
- Multiple Certificate Authentication— Creates an endpoint attribute of type Multiple Certificate Authentication. You can specify the attributes for the multiple certificate authentication of remote VPN users. See Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication, page 32-38.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Add/Edit DAP Entry Dialog Box > AAA Attributes Cisco

To configure AAA attributes as selection criteria for dynamic access policies, in the Add/Edit DAP Entry dialog box, set AAA Attributes Cisco as the selection criterion to be used to select and apply the dynamic access policies during session establishment. You can set these attributes either to match or not match the value you enter. There is no limit for the number of AAA attributes for each dynamic access policy.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **AAA Attributes Cisco** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7

- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-6 Add/Edit DAP Entry Dialog Box > AAA Attributes Cisco

Element	Description
Criterion	Shows AAA Attributes Cisco as the selection criterion.
Group Policy	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the name of the AAA server group associated with the user. The maximum length is 64 characters. AAA server groups represent collections of authentication servers focused on enforcing specific aspects of your overall network security policy.
IPv4 Address	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the assigned IP address. Addresses are predefined network objects. You can also click **Select** to open a dialog box that lists all available network hosts, and in which you can create or edit network host objects. **Tip** If you select this option and later look at the rule in ASDM, the IP Address attribute is called Assigned IP Address.
IPv6 Address (Security Manager version 4.12 or later and ASA version 9.0 or later)	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the assigned IP address. Addresses are predefined network objects. You can also click **Select** to open a dialog box that lists all available network hosts, and in which you can create or edit network host objects. **Tip** If you select this option and later look at the rule in ASDM, the IP Address attribute is called Assigned IP Address.
Member-of	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter a comma-separated string of group policy names that apply to the user. This attribute lets you indicate multiple group membership. The maximum length is 128 characters. **Tip** If you select this option, and later look at the rule in ASDM, this option will not appear. In general, this option is not used because it can be confused with the memberof LDAP attribute. Because this rule applies to Local authentication, you might want to use the Username attribute instead of the Member-of attribute.
Username	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the username of the authenticated user. A maximum of 64 characters is allowed.
Username 2	Select the check box, select the matching criteria (*is* or *isn't*) from the drop-down list, and enter the secondary username of the authenticated user.

Table 32-6 Add/Edit DAP Entry Dialog Box > AAA Attributes Cisco (continued)

Element	Description
Connection Profiles	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and select the connection profile from a list of all the SSL VPN Connection Profile policies defined on the security appliance. An SSL VPN connection profile comprises a set of records that contain VPN tunnel connection profile policies, including the attributes that pertain to creating the tunnel itself. **Note** For a description of the procedure to configure an SSL VPN Connection Profiles policy, see Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7.
SCEP Required	Select the check box, select the matching criteria (*is* or *isn't*) from the drop-down list, and select *True* or *False*. This attribute enables to match whether or not the connection fails the certificate authentication.

Add/Edit DAP Entry Dialog Box > AAA Attributes LDAP

The LDAP client stores all native LDAP response attribute value pairs in a database associated with the AAA session for the user. The LDAP client writes the response attributes to the database in the order in which it receives them. It discards all subsequent attributes with that name. This scenario might occur when a user record and a group record are both read from the LDAP server. The user record attributes are read first, and always have priority over group record attributes.

To support Active Directory group membership, the AAA LDAP client provides special handling of the LDAP memberOf response attribute. The AD memberOf attribute specifies the DN string of a group record in AD. The name of the group is the first CN value in the DN string. The LDAP client extracts the group name from the DN string and stores it as the AAA memberOf attribute, and in the response attribute database as the LDAP memberOf attribute. If there are additional memberOf attributes in the LDAP response message, then the group name is extracted from those attributes and is combined with the earlier AAA memberOf attribute to form a comma separated string of group names, also updated in the response attribute database.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **AAA Attributes LDAP** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-7 Add/Edit DAP Entry Dialog Box > AAA Attributes LDAP

Element	Description
Criterion	Shows AAA Attributes LDAP as the selection criterion.
Attribute ID	Specify the name of the LDAP attribute map in the dynamic access policy. LDAP attribute maps take the attribute names that you define and map them to Cisco-defined attributes. A maximum of 64 characters is allowed.
Value	Select the matching criteria (for example, *is*) from the drop-down list, and enter the custom map value that maps to a Cisco Map Value or enter the Cisco map value that maps to the Custom Map Value. To enter multiple values, separate each value with **;** as the delimiter. The attribute map is populated with value mappings that apply customer, user-defined attribute values to the customer attribute name and to the matching Cisco attribute name and value. Alternatively, click the **Fetch AD Groups** button to open the Fetch AD Groups dialog box. The table in the dialog box lists the UserGroup ID and UserGroup Name of the available LDAP servers that you can choose from. Select one or more rows and click the **Select** button. To search for a particular UserGroup in the list you can enter text in the Filter text box and click **Search**. The UserGroup name meeting the criteria appears in the list. **Note** To be able to view the list of available LDAP servers you must first configure the mapping of Domain to AD Server Group. To perform this task, go to **Tools >Security Manager Administration** and select **Identity Settings** from the table of contents. For more information, see Identity Settings Page, page 11-38.

Add/Edit DAP Entry Dialog Box > AAA Attributes RADIUS

The RADIUS client stores all native RADIUS response attribute value pairs in a database associated with the AAA session for the user. The RADIUS client writes the response attributes to the database in the order in which it receives them. It discards all subsequent attributes with that name. This scenario might occur when a user record and a group record are both read from the RADIUS server. The user record attributes are read first, and always have priority over group record attributes.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **AAA Attributes RADIUS** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-8 Add/Edit DAP Entry Dialog Box > AAA Attributes RADIUS

Element	Description
Criterion	Shows AAA Attributes RADIUS as the selection criterion.
Attribute ID	Specify the name of the RADIUS attribute name or number in the dynamic access policy. A maximum of 64 characters is allowed. RADIUS attribute names do not contain the cVPN3000 prefix to better reflect support for all three security appliances (VPN 3000, PIX, and the ASA). The appliances enforce the RADIUS attributes based on attribute numeric ID, not attribute name. LDAP attributes are enforced by their name, not by the ID.
Value	Select the matching criteria (for example, *is*) from the drop-down list, and enter the attribute value.

Add/Edit DAP Entry Dialog Box > Anti-Spyware

You can use the Host Scan feature of the Cisco Secure Desktop feature to enable Endpoint Assessment, a scan for antivirus, personal firewall, and antispyware applications and updates that are running on the remote computer. Following the configuration of the prelogin policies and host scan options, you can configure a match of any one or any combination of the Host Scan results to assign a dynamic access policy following the user login.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Anti-Spyware** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-9 Add/Edit DAP Entry Dialog Box > Anti-Spyware

Element	Description
Criterion	Shows Anti-Spyware as the selection criterion.
Type	Select one of the following options and assign the associated values: • **Not Installed**—Select if the absence of the named anti-spyware from the remote PC is sufficient to match the prelogin policy you are configuring. • **Installed and enabled**—Select if the named anti-spyware must be present and enabled on the remote PC to match the prelogin policy you are configuring. • **Installed and disabled**—Select if the mere presence of the named anti-spyware on the remote PC is sufficient to match the prelogin policy you are configuring.
Vendor Name	Select the text that describes the application vendor from the list.
Product ID	Select a unique identifier for the product that is supported by the selected vendor from the list.
Product Description	Available only if you selected **Matches** as the Type. Select the check box, then select the description of the product from the list.
Version	Available only if you selected **Matches** as the Type. Identify the version of the application, and specify whether you want the endpoint attribute to be equal to/not equal to that version.
Last Update	Available only if you selected **Matches** as the Type. Specify the number of days since the last update. You might want to indicate that an update should occur in less than or greater than the number of days you enter here.

Add/Edit DAP Entry Dialog Box > Anti-Virus

You can configure a scan for antivirus applications and updates as a condition for the completion of a Cisco AnyConnect or clientless SSL VPN connection. Following the prelogin assessment, Cisco Secure Desktop loads Endpoint Assessment checks and reports the results back to the security appliance for use in assigning a dynamic access policy.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Anti-Virus** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-10 Add/Edit DAP Entry Dialog Box > Anti-Virus

Element	Description
Criterion	Shows Anti-Virus as the selection criterion.
Type	Select one of the following options and assign the associated values: • **Not Installed**—Select if the absence of the named anti-virus from the remote PC is sufficient to match the prelogin policy you are configuring. • **Installed and enabled**—Select if the named anti-virus must be present and enabled on the remote PC to match the prelogin policy you are configuring. • **Installed and disabled**—Select if the mere presence of the named anti-virus on the remote PC is sufficient to match the prelogin policy you are configuring.
Vendor Name	Select the text that describes the application vendor from the list.
Product ID	Select a unique identifier for the product that is supported by the selected vendor from the list.
Product Description	Available only if you selected the criteria to match the endpoint attribute for the dynamic access policy. Select the check box, then select the description of the product from the list.
Version	Available only if you selected the criteria to match the endpoint attribute for the dynamic access policy. Identify the version of the application, and specify whether you want the endpoint attribute to be equal to/not equal to that version.
Last Update	Available only if you selected the criteria to match the endpoint attribute for the dynamic access policy. Specify the number of days since the last update. You might want to indicate that an update should occur in less than or greater than the number of days you enter here.

Add/Edit DAP Entry Dialog Box > AnyConnect Identity

To configure AnyConnect Identity attributes as selection criteria for dynamic access policies, set AnyConnect Identity as the selection criterion in the Add/Edit DAP Entry dialog box. The ASA generates DAP endpoint attributes based on the AnyConnect Identification attributes received from the AnyConnect mobile client. You are not required to enable Cisco Secure Desktop to configure these specific attributes using Security Manager.

For the purposes of assigning a dynamic access policy, if you configure more than one AnyConnect Identity attribute for a particular DAP entry, the entry will be considered a match if any of the attributes values are true. There is no limit for the number of AnyConnect Identity attributes for each dynamic access policy.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **AnyConnect Identity** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-11 Add/Edit DAP Entry Dialog Box > AnyConnect Identity

Element	Description
Criterion	Shows AnyConnect Identity as the selection criterion.
Client Version	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the AnyConnect Client version number.
Platform	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and select the appropriate platform from the drop-down list.
Platform Version	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the appropriate version number of the platform.
Device Type	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and select the appropriate device type from the drop-down list.
Device Unique ID	Select the check box, select the matching criteria (for example, *is*) from the drop-down list, and enter the unique device ID. This ID distinguishes the device allowing you to set policies exclusive to that device.

Add/Edit DAP Entry Dialog Box > Application

Use this dialog box to indicate the type of remote access connection as the endpoint attribute for the dynamic access policy.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Application** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-12 Add/Edit DAP Entry Dialog Box > Application

Element	Description
Criterion	Shows Application as the selection criterion.
Client Type	Select the check box, then select the matching criteria (for example, *is* or *isn't*) from the drop-down list, and specify the type of remote access connection from the list: AnyConnect, Clientless, Cut-through Proxy, IPsec, Generic IKEv2 Client, or L2TP. **Note** If you select AnyConnect as the client type, make sure to enable Cisco Secure Desktop. If it is not enabled, Security Manager generates an error.

Add/Edit DAP Entry Dialog Box > Device

The DAP Device Criterion lets you provide specific device information for use during the associated prelogin policy checking. You can provide one or more of the following attributes for a device—host name, MAC address, port number, Privacy Protection selection—and indicate whether each *is* or *isn't* to be matched.

Note that *isn't* is exclusionary. For example, if you specify the criterion `Host Name isn't zulu_2`, all devices not named `zulu_2` will match.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Choose **Device** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-13 Add/Edit DAP Entry Dialog Box > Device

Element	Description
Criterion	Shows Device as the selected Criterion.
Host Name	Select this option, choose a match criterion (*is* or *isn't*) from the related drop-down list, and then enter the device host name to be matched.
MAC Address	Select this option, choose a match criterion (*is* or *isn't*) from the related drop-down list, and then enter the device's MAC address to be matched.
BIOS Serial Number	Select this option, choose a match criterion (*is* or *isn't*) from the related drop-down list, and then enter the BIOS serial number value of the device you are matching for. The number format is manufacturer-specific. There is no format requirement.
Port Number	Select this option, choose a match criterion (*is* or *isn't*), and then enter or Select the device port to be matched.
TCP/UDP Port Number	Select this option, choose a match criterion (*is* or *isn't*), and then enter or Select the TCP/UDP port in listening state that you are matching for. In the TCP/UDP combo box, select the kind of port you are matching for: TCP (IPv4), UDP(IPv4), TCP(IPv6) or UDP(IPv6). Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running the version 9.0 or later. If you are matching for more than one port, make several individual endpoint attribute rules in the DAP and specify one port in each.
Privacy Protection	Select this option, choose a match criterion (*is* or *isn't*), and then choose the Privacy Protection option defined on the device: **none**, **cache cleaner**, or **secure desktop**.
CSD Version	Select this option, choose a match criterion (*is* or *isn't*) from the related drop-down list, and then enter the version of the Host Scan image running on the endpoint.
Endpoint Assessment Version	Select this option, choose a match criterion (*is* or *isn't*) from the related drop-down list, and then enter the version of endpoint assessment (OPSWAT) you are matching for.

Add/Edit DAP Entry Dialog Box > File

The file criterion prelogin check lets you specify that a certain file must or must not exist to be eligible for the associated prelogin policy. For example, you might want to use a file prelogin check to ensure a corporate file is present or one or more peer-to-peer file-sharing programs containing malware are not present before assigning a prelogin policy.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **File** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-14 Add/Edit DAP Entry Dialog Box > File

Element	Description
Criterion	Shows File as the selection criterion.
Type	Specify whether this endpoint attribute must match or not match the criteria configured for selecting and applying dynamic access policies during session establishment.
Endpoint ID	Select a string that identifies an endpoint for files. Dynamic access policies use this ID to match Cisco Secure Desktop host scan attributes for dynamic access policy selection. You must configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it, reducing the possibility of errors in typing or syntax.
Filename	Specify the filename.
Last Update	Available only if you selected the criteria to match the endpoint attribute for the dynamic access policy. Specify the number of days since the last update. You might want to indicate that an update should occur in less than (<) or more than (>) the number of days you enter here.
Checksum	Available only if you selected the criteria to match the endpoint attribute for the DAP record. Select the check box to specify a checksum to authenticate the file, then enter a checksum in hexadecimal format, beginning with 0x. Beginning with version 4.7, Security Manager provides a utility to compute CRC32 checksum for a file. Click the **Compute CRC32 Checksum** button to open the Compute Checksum dialog box. Click Browse to open the File browser, select the required file and then click the Compute button. The CRC32 checksum of the file will be calculated and populated in the Checksum field. **Note** Only client-side browsing is supported for the Compute CRC32 Checksum utility. By default client-side browsing is enabled. If you have disabled it, you must enable it by selecting **Tools** > **Security Manager Administration** and select **Customize Desktop** from the table of contents. For more information, see Customize Desktop Page, page 11-10.

Add/Edit DAP Entry Dialog Box > NAC

NAC protects the enterprise network from intrusion and infection from worms, viruses, and rogue applications by performing endpoint compliancy and vulnerability checks as a condition for production access to the network. We refer to these checks as *posture†validation*. You can configure posture

validation to ensure that the anti-virus files, personal firewall rules, or intrusion protection software on a host with an AnyConnect or Clientless SSL VPN session are up-to-date before providing access to vulnerable hosts on the intranet. Posture validation can include the verification that the applications running on the remote hosts are updated with the latest patches. NAC occurs only after user authentication and the setup of the tunnel. NAC is especially useful for protecting the enterprise network from hosts that are not subject to automatic network policy enforcement, such as home PCs. The security appliance uses Extensible Authentication Protocol (EAP) over UDP (EAPoUDP) messaging to validate the posture of remote hosts.

The establishment of a tunnel between the endpoint and the security appliance triggers posture validation. You can configure the security appliance to pass the IP address of the client to an optional audit server if the client does not respond to a posture validation request. The audit server, such as a Trend server, uses the host IP address to challenge the host directly to assess its health. For example, it may challenge the host to determine whether its virus checking software is active and up-to-date. After the audit server completes its interaction with the remote host, it passes a token to the posture validation server, indicating the health of the remote host.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **NAC** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-15 Add/Edit DAP Entry Dialog Box > NAC

Element	Description
Criterion	Shows NAC as the selection criterion.
Posture Status	Select the matching criteria (for example, *is*) from the drop-down list, then enter the posture token string received from ACS.

Add/Edit DAP Entry Dialog Box > Operating System

The prelogin assessment includes a check for the OS attempting to establish a VPN connection. When the user attempts to connect, however, Cisco Secure Desktop checks for the OS, regardless of whether you insert an OS prelogin check.

If the prelogin policy assigned to the connection has Secure Desktop (Secure Session) enabled and if the remote PC is running Microsoft Windows XP or Windows 2000, it installs Secure Session, regardless of whether you insert an OS prelogin check. If the prelogin policy has Secure Desktop enabled and the operating system is Microsoft Windows Vista, Mac OS X 10.4, or Linux, Cache Cleaner runs instead. Therefore, you should make sure the Cache Cleaner settings are appropriate for a prelogin policy on

which you have configured Secure Desktop or Cache Cleaner to install. Although Cisco Secure Desktop checks for the OS, you may want to insert an OS prelogin check as a condition for applying a prelogin policy to isolate subsequent checks for each OS.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Operating System** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-16 Add/Edit DAP Entry Dialog Box > Operating System

Element	Description
Criterion	Shows Operating System as the selection criterion.
OS Version	Select the check box, then select the matching criteria (for example, *is*) from the drop-down list, and select the OS version from the list. Select Apple Plugin for iPhones and similar devices.
Service Pack	Select the check box, then select the matching criteria (for example, *is*) from the drop-down list, and select the service pack for the operating system.

Add/Edit DAP Entry Dialog Box > Personal Firewall

You can click Host Scan in the Cisco Secure Desktop interface to enable Endpoint Assessment, a scan for personal firewalls that are running on the remote computer. Most, but not all, personal firewall programs support active scan, which means that the programs are memory-resident, and therefore always running.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **AAA Attributes Cisco** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4

- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-17 Add/Edit DAP Entry Dialog Box > Personal Firewall

Element	Description
Criterion	Shows Personal Firewall as the selection criterion.
Type	Select one of the following options and assign the associated values: • **Not Installed**—Select if the absence of the named personal firewall from the remote PC is sufficient to match the prelogin policy you are configuring. • **Installed and enabled**—Select if the named personal firewall must be present and enabled on the remote PC to match the prelogin policy you are configuring. • **Installed and disabled**—Select if the mere presence of the named personal firewall on the remote PC is sufficient to match the prelogin policy you are configuring.
Vendor Name	Select the text that describes the application vendor from the list.
Product ID	Select a unique identifier for the product that is supported by the selected vendor from the list.
Product Description	Available only if you selected that this endpoint attribute and all its settings must be available on the remote PC. Select the check box, then select the description of the product from the list.
Version	Available only if you selected that this endpoint attribute and all its settings must be available on the remote PC. Identify the version of the application, and specify whether you want the endpoint attribute to be equal to/not equal to that version.

Add/Edit DAP Entry Dialog Box > Policy

Windows locations let you determine how clients connect to your virtual private network, and protect it accordingly. For example, clients connecting from within a workplace LAN on a 10.x.x.x network behind a NAT device are an unlikely risk for exposing confidential information. For these clients, you might set up a Cisco Secure Desktop Windows Location named Work that is specified by IP addresses on the 10.x.x.x network, and disable both the Cache Cleaner and the Secure Desktop function for this location. Cisco Secure Desktop checks locations in the order listed on the Windows Location Settings window, and grants privileges to client PCs based on the first location definition they match.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Policy** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-18 Add/Edit DAP Entry Dialog Box > Policy

Element	Description
Criterion	Shows Policy as the selection criterion.
Location	Select the matching criteria (for example, *is*) from the drop-down list, and select the Cisco Secure Desktop Microsoft Windows location profile from the list. All the locations configured in the Cisco Secure Desktop Manager are displayed in this list.

Add/Edit DAP Entry Dialog Box > Process

You can specify a set of process names, which form a part of Basic Host Scan. The host scan, which includes Basic Host Scan and Endpoint Assessment, or Advanced Endpoint Assessment; occurs after the prelogin assessment but before the assignment of a dynamic access policy. Following the Basic Host Scan, the security appliance uses the login credentials, the host scan results, prelogin policy, and other criteria you configure to assign a DAP.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Process** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-19 Add/Edit DAP Entry Dialog Box > Process

Element	Description
Criterion	Shows Process as the selection criterion.

Table 32-19 Add/Edit DAP Entry Dialog Box > Process (continued)

Element	Description
Type	Select one of the following options and assign the associated values: • **Matches**—Select if the mere presence of the named process on the remote PC is sufficient to match the prelogin policy you are configuring. • **Doesn't Match**—Select if the absence of the named process from the remote PC is sufficient to match the prelogin policy you are configuring.
Endpoint ID	A string that identifies an endpoint for files, processes or registry entries. Dynamic access policies use this ID to match Cisco Secure Desktop host scan attributes for dynamic access policy selection. You must configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it, reducing the possibility of errors in typing or syntax.
Path	Select the check box, then select the matching criteria (for example, *is*) from the drop-down list, and enter the name of the process. You can display it in Microsoft Windows by opening the Windows Task Manager window and clicking the Processes tab. Configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it and specify the same index when you assign this entry as an endpoint attribute when configuring a DAP, reducing the possibility of errors in typing or syntax.

Add/Edit DAP Entry Dialog Box > Registry

Registry key scans apply only to computers running Windows Microsoft Windows operating systems. Basic Host Scan ignores registry key scans if the computer is running Mac OS or Linux.

Note Duplicate entries are not allowed. If you configure a dynamic access policy with no AAA or endpoint attributes, the security appliance always selects it since all selection criteria are satisfied.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Registry** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-20 Add/Edit DAP Entry Dialog Box > Registry

Element	Description
Criterion	Shows Registry as the selection criterion.
Type	Select one of the following options and assign the associated values: • **Matches**—Select if the mere presence of the named registry key on the remote PC is sufficient to match the prelogin policy you are configuring. For example, select this option if you want to require the following registry key to be present to match a criterion for assigning a prelogin policy: HKEY_LOCAL_MACHINE\SOFTWARE\<Protective_Software > • **Doesn't Match**—Select if the absence of the named registry key from the remote PC is sufficient to match the prelogin policy you are configuring. For example, select this option if you want to require the following registry key to be absent to match a criterion for assigning a prelogin policy: HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Run\<Evil_SpyWare>
Endpoint ID	A string that identifies an endpoint for files, processes or registry entries. Dynamic access policies use this ID to match Cisco Secure Desktop host scan attributes for dynamic access policy selection. You must configure Host Scan before you configure this attribute. When you configure Host Scan, the configuration displays in this pane, so you can select it, reducing the possibility of errors in typing or syntax.
Registry Name	Select the text that describes the registry name from the list.
Value	Select the value, **dword** or **string**, from the list, then select the matching criteria (whether it equals or does not equal), and enter a decimal or a string to compare with the dword or string value of the registry key on the remote PC. **Note** "DWORD" refers to the attribute in the Add/Edit Registry Criterion dialog box. "Dword" refers to the attribute as it appears in the registry key. Use the regedit application, accessed on the Windows command line, to view the Dword value of a registry key, or use it to add a Dword value to the registry key to satisfy the requirement you are configuring.
Ignore Case	When selected, ignores the case in the registry entry if it includes a string.

Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication

The DAP multiple certificate authentication criterion allows you to provide specific certificate information for use during the associated prelogin policy checking. Cisco Security Manager supports two certificates to authenticate remote VPN users. You can provide one or more of the following attributes for the certificates—subject, issuer, subject alternate name, serial number and certificate store.

Note You can modify the DAP entry except the certificate option.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12 with the Main tab selected, then click **Create**, or select a dynamic access policy in the table and click **Edit**. The Add/Edit DAP Entry dialog box is displayed. Select **Multiple Certificate Authentication** as the Criterion.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-21 Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication

Element	Description
Criterion	Shows Multiple Certificate Authentication as the selection criterion.
Certificate	Multiple certification in 4.13 refers to two certificate authentication. Select one of the following options and assign the associated attributes: • **Cert1**—Select to provide the certificate 1 details to match the prelogin policy you are configuring. • **Cert2**—Select to provide the certificate 2 details to match the prelogin policy you are configuring. **Note** You cannot edit/modify the certificate option.

Table 32-21 Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication (continued)

Element	Description
Subject	From the drop-down list, select the domain name (DN) attribute field from subject name of the certificate: • dnq—Domain name qualifier • fulldn—Full subject-name • ser—Serial number • cn—Common name • i—Initials • ou—Organization Unit • sp—State/Province • o—Organization • n—Name • sn—Surname • t—Title • uid—User Identifier • genq—Generation Qualifier • c—Country • l—Locality • gn—Given Name • ea—E-mail address In the adjacent text box, enter the DAP entry value for the selected Subject. **Note** If you leave the text box blank, and error message appears while saving.

Table 32-21 Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication (continued)

Element	Description
Issuer	From the drop-down list, select the domain name (DN) attribute field from issuer name of the certificate: • dnq—Domain name qualifier • fulldn—Full issuer-name • ser—Serial number • cn—Common name • i—Initials • ou—Organization Unit • sp—State/Province • o—Organization • n—Name • sn—Surname • t—Title • uid—User Identifier • genq—Generation Qualifier • c—Country • l—Locality • gn—Given Name • ea—E-mail address In the adjacent text box, enter the DAP entry value for the selected issuer. **Note** If you leave the text box blank, an error message appears while saving.
Subject Alternate Name	For configuring the serial number, select upn from this drop-down list. In the adjacent text box, enter the User Principal Name from Subject Alt Name field of certificate
Serial Number	Enter the serial number of certificate to be matched. This value should be a hexadecimal number (a combination of 0 to 9, and A to F). **Note** If you enter a non-hexadecimal number, an error message appears while saving.

Table 32-21 Add/Edit DAP Entry Dialog Box > Multiple Certificate Authentication (continued)

Element	Description
Certificate Store	Select the relevant store from where the certificate can be found for authentication: • None—Choose if you are not aware of the certificate type. • Machine—Choose if the certificate pertains to machine (accessible only by privileged processes). You cannot select this option for both cert1 and cert2. • User—Choose if the certificate pertains to user log in (accessible only by processes owned by the logged-in user). **Note** For Windows, the store may be a) one machine and one user, or b) two users. For non-Windows platforms, the indication is always two user certificates.

Logical Operations Tab

Use the Logical Operations tab of the Add/Edit Dynamic Access Policy dialog box to configure multiple instances of the AAA and each type of endpoint attribute that you defined in the DAP Entry dialog box. On this tab, set each type of endpoint or AAA attribute to require only one instance of a type (Match Any = OR) or to have all instances of a type (Match All = AND).

- If you configure only one instance of an endpoint category, you do not need to set a value.
- For some endpoint attributes, it is not useful to configure multiple instances. For example, no users have more than one running OS.
- You are configuring the Match Any/Match All operation within each endpoint type. The security appliance evaluates each type of endpoint attribute, and then performs a logical AND operation on all of the configured endpoints. That is, each user must satisfy the conditions of ALL of the endpoints you configure, as well as the AAA attributes.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12, then click the **Logical Operations** tab.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-22 Add/Edit Dynamic Access Policy Dialog Box > Logical Operations Tab

Element	Description
AAA	Select one of the following options if you defined the AAA attribute in the dynamic access policy: • **Match Any**—Creates an OR relationship among the attributes. Attributes matching any of your criteria are included in the filter. The security appliance grants access to a particular user for a particular session even if any one of the attributes is matching all your criteria. • **Match All**—Creates an AND relationship among the attributes. The security appliance grants access to a particular user for a particular session only if the attributes are matching all your criteria. • **Match None**—Creates a NOT relationship among the attributes. The dynamic access policy specifies that none of the attributes of the user need to match to be granted access to a session.
Anti-Spyware	Select one of the following options if you defined Anti-Spyware as an endpoint attribute: • **Match Any**—Creates an OR relationship among the attributes. Policies matching any instance of your criteria are used to authorize users. • **Match All**—Creates an AND relationship among the attributes. Only those attributes matching all your criteria are used to authorize users.
Anti-Virus	Select one of the following options if you defined Anti-Virus as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.
Application	Select one of the following options if you defined Application as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.

Table 32-22 Add/Edit Dynamic Access Policy Dialog Box > Logical Operations Tab (continued)

Element	Description
File	Select one of the following options if you defined File as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.
Personal Firewall	Personal firewall rules let you specify applications and ports for the firewall to allow or block. Select one of the following options if you defined Personal Firewall as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.
Process	Select one of the following options if you defined Process as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.
Registry	Registry key scans apply only to computers running Windows Microsoft Windows operating systems. Basic Host Scan ignores registry key scans if the computer is running Mac OS or Linux. Select one of the following options if you defined Registry as an endpoint attribute: • **Match Any**—Set to require that user authorization attributes match any of the values in the Antivirus endpoint attributes you are configuring. • **Match All**—Set to require that user authorization attributes match all of the values in the endpoint attributes you are configuring, as well as satisfying the AAA attribute.

Advanced Expressions Tab

Use the Advanced Expressions tab of the Add/Edit Dynamic Access Policy dialog box to set additional attributes for the dynamic access policy. You can configure multiple instances of each type of endpoint attribute. Be aware that this is an advanced feature that requires knowledge of LUA (www.lua.org).

Note For detailed information about advanced expressions, see About Advanced Expressions for AAA or Endpoint Attributes and Examples of DAP Logical Expressions.

Navigation Path

Open the Add/Edit Dynamic Access Policy Dialog Box, page 32-12, then click the **Advanced Expressions** tab.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Field Reference

Table 32-23 Add/Edit Dynamic Access Policy Dialog Box > Advanced Expressions Tab

Element	Description
Basic Expressions	This text box is populated with basic expressions based on the endpoint and AAA attributes that you configured in the dynamic access policy.
Relationship Drop-down List	Specify the relationship between the basic selection rules and the logical expressions you enter on this tab, that is, whether the new attributes add to or substitute for the AAA and endpoint attributes already set. Select one of the following options: • **Basic AND Advanced**—Creates an AND relationship between the basic and advanced expressions. Both the basic and advanced expressions defined in the dynamic access policy are considered while authenticating users. By default, this option is selected. • **Basic OR Advanced**—Creates an OR relationship between the basic and advanced expressions. Users are granted access to a session if either the basic or advanced expressions in the dynamic access policy are matched with the user policy. • **Basic Only**—Only the basic expressions defined in the DAP entry are used to determine whether the security appliance grants users access to a particular session. • **Advanced Only**—Only the advanced expressions defined in the DAP entry are used to authorize users for an SSL VPN session.
Advanced Expressions	Enter one or more logical expressions to set AAA or endpoint attributes other than what is possible in the AAA and Endpoint areas above. Enter free-form LUA text that defines new AAA and/or endpoint selection attributes. Security Manager does not validate text that you enter here; it just copies this text to the dynamic access policy XML file, and the security appliance processes it, discarding any expressions it cannot parse.

Cisco Secure Desktop Manager Policy Editor Dialog Box

Using the Cisco Secure Desktop Manager (CSDM) Policy Editor dialog box, you can configure prelogin policies, specify the checks to be performed between the time the user establishes a connection with the security appliance and the time the user enters the login credentials, and configure host scans. For an explanation of configuring CSD on an ASA device, see Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9.

Note The Cisco Secure Desktop Manager Policy Editor is an independent program. For information about configuring CSD, and what CSD can do for you, see the materials available online at http://www.cisco.com/en/US/products/ps6742/tsd_products_support_configure.html. Look specifically for information on configuring prelogin policies and host scan. Select the configuration guide for the CSD version you are configuring.

Navigation Path

Open the Dynamic Access Page (ASA), page 32-11, then click **Configure** from the Cisco Secure Desktop section (you must first specify a CSD package). The CSDM Policy Editor dialog box is displayed.

Related Topics

- Understanding DAP Attributes, page 32-4
- Configuring DAP Attributes, page 32-7
- Configuring Dynamic Access Policies, page 32-2

Managing Remote Access VPNs on IOS and PIX 6.3 Devices

Note From version 4.17, though Cisco Security Manager continues to support IOS and PIX features/functionality, it does not support any enhancements.

You can configure and manage remote access IPsec on devices running Cisco IOS Software or PIX 6.3, and SSL VPNs on IOS 12.4(6)T or higher devices (but not on PIX devices). For more information on the specific device models supported, see Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8.

The configuration of these remote access VPNs are the same for these device types. ASA and PIX 7.0+ devices use different configurations for remote access VPNs (as explained in Chapter 31, "Managing Remote Access VPNs on ASA and PIX 7.0+ Devices").

The topics in this chapter explain how to configure policies that are specific to IOS and PIX 6.3 devices. Additionally, review the following topics for more information about remote access VPNs:

- Understanding Remote Access VPNs, page 30-1
- Understanding Devices Supported by Each Remote Access VPN Technology, page 30-8
- Discovering Remote Access VPN Policies, page 30-12
- Using the Remote Access VPN Configuration Wizard, page 30-13
  - Creating IPSec VPNs Using the Remote Access VPN Configuration Wizard (IOS and PIX 6.3 Devices), page 30-36
  - Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32

This chapter contains the following topics:

- Overview of Remote Access VPN Policies for IOS and PIX 6.3 Devices, page 33-2
- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3
- Configuring High Availability in Remote Access VPNs (IOS), page 33-11
- Configuring User Group Policies, page 33-13
- Configuring an SSL VPN Policy (IOS), page 33-14

Overview of Remote Access VPN Policies for IOS and PIX 6.3 Devices

Note From version 4.17, though Cisco Security Manager continues to support IOS and PIX features/functionality, it does not support any enhancements.

When you configure remote access VPNs on IOS or PIX 6.3 devices, you use the following policies based on the type of VPN you are configuring. Note that you cannot configure SSL VPNs on PIX 6.3 devices.

- **Policies used with both IPsec and SSL remote access VPNs:**
  - **Global Settings**—You can define global settings that apply to all devices in your remote access VPNs. These settings include Internet Key Exchange (IKE), IPsec, NAT, and fragmentation definitions. The global settings typically have defaults that work in most situations, so configuring the Global Settings policy is optional; configure it only if you need non-default behavior. For more information, see Configuring VPN Global Settings, page 26-30.
  - **Public Key Infrastructure**—You can create a Public Key Infrastructure (PKI) policy to generate enrollment requests for CA certificates and RSA keys, and to manage keys and certificates. Certification Authority (CA) servers are used to manage these certificate requests and issue certificates to users who connect to your IPsec or SSL remote access VPN. For more information, see Understanding Public Key Infrastructure Policies, page 26-51 and Configuring Public Key Infrastructure Policies for Remote Access VPNs, page 26-56.
- **Policies used in remote access IPsec VPNs only:**
  - **IKE Proposal**—Internet Key Exchange (IKE), also called ISAKMP, is the negotiation protocol that enables two hosts to agree on how to build an IPsec security association. IKE is used to authenticate IPsec peers, negotiate and distribute IPsec encryption keys, and to automatically establish IPsec security associations (SAs). Use the IKE Proposal policy to define the requirements for phase 1 of the IKE negotiation. For more information, see Configuring an IKE Proposal, page 26-9.
  - **IPsec Proposal (IOS/PIX 6.x)**—An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA. The policy is used for IKE phase 2 negotiations. For more information, see Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3.
  - **High Availability**—High Availability (HA) is supported by the creation of an HA group made up of two or more hub devices that use Hot Standby Routing Protocol (HSRP) to provide transparent, automatic device failover. For more information, see Configuring High Availability in Remote Access VPNs (IOS), page 33-11.
  - **User Groups (IOS/PIX 6.x)**—A user group policy specifies the attributes that determine user access to and use of the VPN. For more information, see Configuring User Group Policies, page 33-13.
- **Policies used in remote access SSL VPNs only:**
  - **SSL VPN**—The SSL VPN policy table lists all of the contexts that define the virtual configurations of the SSL VPN. Each context has a gateway, domain or virtual hostname, and user group policies. For more information, see Configuring an SSL VPN Policy (IOS), page 33-14.

Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices)

Note From version 4.17, though Cisco Security Manager continues to support IOS and PIX features/functionality, it does not support any enhancements.

This procedure describes how to create or edit an IPsec proposal for your remote access VPN server when the server uses Cisco IOS Software or PIX release 6.3.

An IPsec proposal is a collection of one or more crypto maps. A crypto map combines all the components required to set up IPsec security associations (SAs), including IPsec rules, transform sets, remote peers, and other parameters that might be necessary to define an IPsec SA.

When configuring an IPsec proposal, you must define the external interface through which the remote access clients connect to the server, and the encryption and authentication algorithms that protect the data in the VPN tunnel. You can also select a group authorization (Group Policy Lookup) method that defines the order in which group policies are searched (on the local server or on external AAA servers) and a user authentication (Xauth) method that defines the order in which user accounts are searched.

For more information on IPsec tunnel concepts, see Understanding IPsec Proposals, page 26-18.

When you create or edit an IPsec proposal, you can also configure:

- A VPN Services Module (VPNSM) interface IPsec VPN Shared Port Adapter (VPN SPA) on a Catalyst 6500/7600 device (see VPNSM/VPN SPA/VSPA Settings Dialog Box, page 33-6).
- A dynamic virtual interface on an IOS router running Cisco IOS Software version 12.4(2)T or later, except 7600 device. For more information, see Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices), page 33-7.
- VRF-Aware IPsec on a router or Catalyst 6500/7600 device (see Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices), page 33-7).

Related Topics

- Understanding VRF-Aware IPsec, page 25-14
- VPNSM/VPN SPA/VSPA Settings Dialog Box, page 33-6
- Table Columns and Column Heading Features, page 1-49

Step 1 Do one of the following:

- (Device view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (IOS/PIX 6.x)** from the Policy selector.
- (Policy view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (IOS/PIX 6.x)** from the Policy Type selector. Select an existing policy or create a new one.

The IPsec Proposal page opens and lists the configured proposals, including the VPN endpoint, IPsec transform set, and whether reverse route injection is configured for the proposal. You can add other columns to the default display to show the AAA, VRF, and dVTI configuration.

Step 2 Do any of the following:

- To add a new IPsec proposal, click the **Add Row (+)** button and fill in the IPsec Proposal Editor dialog box. For detailed information on the available options, see IPsec Proposal Editor (IOS, PIX 6.3 Devices), page 33-4.

- To edit an existing proposal, select it and click the **Edit Row (pencil)** button.
- To delete a proposal, select it and click the **Delete Row (trash can)** button.

IPsec Proposal Editor (IOS, PIX 6.3 Devices)

Note From version 4.17, though Cisco Security Manager continues to support IOS and PIX features/functionality, it does not support any enhancements.

Use the IPsec Proposal Editor to create or edit an IPsec proposal for an IOS or PIX 6.3 device, including Catalyst 6500/7600, in your remote access VPN. The editor has two tabs—General and Dynamic VTI/VRF Aware IPsec. This topic explains the basic settings on the General tab. For an explanation of Dynamic VTI/VRF Aware IPsec settings, see Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices), page 33-7.

The elements in this dialog box differ according to the selected device. The table below describes the elements on the General tab in the IPsec Proposal Editor dialog box when a Cisco IOS router, Catalyst 6500/7600, or PIX 6.3 device is selected.

Note For a description of the elements in the dialog box when a PIX 7.0+ or ASA device is selected is selected, see IPsec Proposal Editor (ASA, PIX 7.0+ Devices), page 31-41.

Navigation Path

- (Device view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (IOS/PIX 6.x)** from the Policy selector. Click the Add Row (+) or Edit Row (pencil) buttons.
- (Policy view) Select **Remote Access VPN > IPSec VPN > IPsec Proposal (IOS/PIX 6.x)** from the Policy Type selector. Select an existing policy or create a new one. Click the Add Row (+) or Edit Row (pencil) buttons.

Related Topics

- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3
- Understanding IPsec Proposals, page 26-18
- Creating Interface Role Objects, page 6-74
- Creating AAA Server Group Objects, page 6-48

Field Reference

Table 33-1 IPsec Proposal Editor, General Tab, IOS and PIX 6.3 Devices

Element	Description
External Interface	**Note** Available only if the selected device is an IOS router. The external interface through which remote access clients will connect to the server. Enter the name of the interface or interface role object, or click **Select** to select it or to create a new object.

Table 33-1 IPsec Proposal Editor, General Tab, IOS and PIX 6.3 Devices (continued)

Element	Description
Inside VLAN	**Note** Available only if the selected device is a Catalyst 6500/7600. The inside VLAN that serves as the inside interface to the VPN Services Module (VPNSM), VPN SPA, or VSPA. Click **Select** to configure the inside VLAN as explained in VPNSM/VPN SPA/VSPA Settings Dialog Box, page 33-6.
IKEv1 Transform Sets	The transform sets to be used for your tunnel policy. Transform sets specify which authentication and encryption algorithms will be used to secure the traffic in the tunnel. You can select up to nine transform sets. For more information, see Understanding Transform Sets, page 26-20. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security will be used. Click **Select** to select the IPsec transform set policy objects to use in the topology. If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27.
Reverse Route Injection	Reverse Route Injection (RRI) enables static routes to be automatically inserted into the routing process for those networks and hosts protected by a remote tunnel endpoint. For more information, see Understanding Reverse Route Injection, page 26-21. Select one of the following options to configure RRI on the crypto map: • **None**—Disables the configuration of RRI on the crypto map. • **Standard**—Creates routes based on the destination information defined in the crypto map access control list (ACL). This is the default option. • **Remote Peer**—Creates two routes, one for the remote endpoint and one for route recursion to the remote endpoint via the interface to which the crypto map is applied. • **Remote Peer IP**—Specifies an address as the explicit next hop to the remote VPN device. Enter the IP address or a network/host object that specifies the address, or click **Select** to select the network/host object from a list or to create a new object. **Note** If you use network/host objects, you can select the **Allow Value Override per Device** option in the object to override the IP address, if required, for specific devices that use this object.

Table 33-1 IPsec Proposal Editor, General Tab, IOS and PIX 6.3 Devices (continued)

Element	Description
Group Policy Lookup/AAA Authorization Method	The AAA authorization method list that will be used to define the order in which the group policies are searched. Group policies can be configured on both the local server or on an external AAA server. Remote users are grouped, so that when the remote client establishes a successful connection to the VPN server, the group policies for that particular user group are pushed to all clients belonging to the user group. Click **Select** to open a dialog box that lists all available AAA group servers, and in which you can create AAA group server objects. Select all that apply and use the up and down arrow buttons to put them in priority order.
User Authentication (Xauth)/AAA Authentication Method	The AAA or Xauth user authentication method that defines the order in which user accounts are searched. Xauth allows all Cisco IOS software AAA authentication methods to perform user authentication in a separate phase after the IKE authentication phase 1 exchange. Click **Select** to open a dialog box that lists all available AAA group servers, and in which you can create AAA group server objects. Select all that apply and use the up and down arrow buttons to put them in priority order.

VPNSM/VPN SPA/VSPA Settings Dialog Box

Note This dialog box is available only if the selected device is a Catalyst 6500/7600.

Use the VPNSM/VPN SPA/VSPA Settings dialog box to specify the settings for configuring a VPN Services Module (VPNSM), a VPN Shared Port Adapter (VPN SPA), or a Cisco VPN Service Port Adapters (VSPAs) on a Catalyst 6500/7600 device.

Notes

- Before you define the settings, you must import your Catalyst 6500/7600 device to the Security Manager inventory and discover its interfaces. For more information, see Configuring VPNSM or VPN SPA/VSPA Endpoint Settings, page 25-42.
- Before you configure VPNSM or VPN SPA with VRF-Aware IPsec on a device, verify that an IPsec proposal with VRF-Aware IPsec and an IPsec proposal without VRF-Aware IPsec were not configured on the device.

Navigation Path

In the General tab of the IPsec Proposal Editor Dialog Box (for Catalyst 6500/7600 Devices), click **Select** next to the Inside VLAN field. For more information about opening the IPsec Proposal Editor, see IPsec Proposal Editor (IOS, PIX 6.3 Devices), page 33-4.

Related Topics

- Creating Interface Role Objects, page 6-74

Field Reference

Table 33-2 VPNSM/VPN SPA/VSPA Settings Dialog Box

Element	Description
Inside VLAN	The inside VLAN that serves as the inside interface to the VPNSM, VPN SPA, or VSPA, and to which the required crypto maps will be applied. Enter the VLAN ID or click **Select** to select it or to create a new interface role object to identify the VLAN.
Slot Subslot	The number designating the slot location of the VPNSM or VPNSPA/VSPA. If you are configuring a VPNSPA/VSPA, the subslot number is also required. **Note** If you are configuring a VPNSM, select 0.
External Port	The external port or VLAN that connects to the inside VLAN. Enter the name of the VLAN or interface role object, or click **Select** to select it from a list. You must select an interface or interface role that differs from the one selected for the inside VLAN. **Note** If VRF-Aware IPsec is configured on the device, the external port or VLAN must have an IP address. If VRF-Aware IPsec is not configured, the external port or VLAN must not have an IP address.
Enable Failover Blade	Whether to configure a failover VPNSM or VPNSPA/VSPA blade for intra-chassis high availability. **Note** A VPNSM and VPNSPA/VSPA blade cannot be used on the same device as primary and failover blades. Specify the failover blade, as follows: • **Slot**—The slot number that identifies where the VPNSM blade or VPNSPA/VSPA blade is located. • **Subslot**—If you are configuring a VPNSPA/VSPA, select the number of the subslot on which the failover VPN SPA blade is installed. **Note** If you are configuring a VPNSM, select 0.

Configuring Dynamic VTI/VRF Aware IPsec in Remote Access VPNs (IOS Devices)

Note The Dynamic VTI/VRF Aware IPsec tab is available only when the selected device is a Cisco IOS router or Catalyst 6500/7600.

Use the Dynamic VTI/VRF Aware IPsec tab of the IPsec Proposal Editor to configure VRF Aware IPsec settings (on a Cisco IOS router or Catalyst 6500/7600 device), configure a dynamic virtual interface on a Cisco IOS router, or do both, in your remote access VPN.

IOS devices allow dynamic virtual template interfaces (VTIs), which provide highly secure and scalable connectivity for remote-access VPNs, replacing dynamic crypto maps and the dynamic hub-and-spoke method for establishing tunnels. You can use dynamic VTIs for both the server and remote configuration. The tunnels provide an on-demand separate virtual access interface for each VPN session. The configuration of the virtual access interfaces is duplicated from a virtual template configuration, which includes the IPsec configuration and any features configured on the virtual template interface. Dynamic VTIs provide efficiency in the use of IP addresses and provide secure connectivity. They enable dynamically downloadable per-group and per-user policies to be configured on a RADIUS server. Dynamic VTI simplifies VRF-Aware IPsec deployment, as the VRF is configured on the interface.

When this feature is enabled, Security Manager implicitly creates the virtual template interface for the selected device in a remote access VPN. All you must do is provide the IP address on the server that will be used as the virtual template interface, or use an existing loopback interface. The virtual template interface is created on the remote client without an IP address.

Notes

- You can configure dynamic VTI only on routers running Cisco IOS Release 12.4(2)T and later, except 7600 devices.
- You can configure dynamic VTI with or without VRF-Aware IPsec. For more information about VRF-Aware IPsec, see Understanding VRF-Aware IPsec, page 25-14.
- You can also configure dynamic VTI in a site-to-site Easy VPN topology. For more information, see Easy VPN with Dynamic Virtual Tunnel Interfaces, page 28-3.

Navigation Path

In the IPsec Proposal Editor Dialog Box (for IOS routers and Catalyst 6500/7600 devices), click the **Dynamic VTI/VRF Aware IPsec** tab. For more information, see IPsec Proposal Editor (IOS, PIX 6.3 Devices), page 33-4.

Related Topics

- Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3
- Creating Interface Role Objects, page 6-74

Field Reference

Table 33-3 IPsec Proposal Editor, Dynamic VTI/VRF Aware IPsec Tab

Element	Description
Enable Dynamic VTI	When selected, enables Security Manager to implicitly create a dynamic virtual template interface on an IOS router. **Note** Dynamic VTI can be configured only on IOS routers running Cisco IOS Release 12.4(2)T and later, except 7600 devices. If the device does not support Dynamic VTI, the option is greyed out.
Enable VRF Settings	When selected, enables you to configure VRF settings on the device for the selected hub-and-spoke topology. **Note** To remove VRF settings that were defined for the VPN topology, deselect this check box.

Table 33-3 IPsec Proposal Editor, Dynamic VTI/VRF Aware IPsec Tab (continued)

Element	Description
User Group	When you configure a remote access VPN server, remote clients must have the same group name as the user group object configured on the VPN server so that they can connect to the device. Enter the name of the user group policy object associated with the device, or click **Select** to select it from a list. You can also create new objects or edit existing ones from the selection list.
CA Server	Select the Certification Authority (CA) server to use for managing certificate requests for the device. Click **Select** to select the PKI enrollment policy object that defines the CA server, or to create a new object. For more information, see PKI Enrollment Dialog Box, page 26-58. For more information about IPsec configuration with CA servers, see Understanding Public Key Infrastructure Policies, page 26-51.
Virtual Template IP Type	Available if you selected **Enable Dynamic VTI**. Specify the virtual template interface to use: • **IP**—To use an IP address as the virtual template interface. Specify the private IP address. • **Use Loopback Interface**—To use the IP address taken from an existing loopback interface as the virtual template interface. Click **Select** to select the interface or interface role object, or to create a new object that identifies the loopback interface.
VRF Solution	Available if you selected **Enable VRF Settings**. Select the VRF solution: • **1-Box (IPsec Aggregator + MPLS PE)**—One device serves as the Provider Edge (PE) router that does the MPLS tagging of the packets in addition to IPsec encryption and decryption from the Customer Edge (CE) devices. For more information, see VRF-Aware IPsec One-Box Solution, page 25-14. • **2-Box (IPsec Aggregator Only)**—The PE device does only the MPLS tagging, while the IPsec Aggregator device does the IPsec encryption and decryption from the CEs. For more information, see VRF-Aware IPsec Two-Box Solution, page 25-15.
VRF Name	The name of the VRF routing table on the IPsec Aggregator. The VRF name is case-sensitive.

Table 33-3 IPsec Proposal Editor, Dynamic VTI/VRF Aware IPsec Tab (continued)

Element	Description
Route Distinguisher	The unique identifier of the VRF routing table on the IPsec Aggregator. This unique route distinguisher maintains routing separation for each VPN across the MPLS core to the other PE routers. The identifier can be in either of the following formats: • *IP address:X*, where *X* is in the range of 0-999999999. • *N:X*, where *N* is in the range of 0-65535, and X is in the range of 0-999999999. **Note** You cannot override the RD identifier after deploying the VRF configuration to your device. To modify the RD identifier after deployment, you must manually remove it through the device CLI and then deploy again.
Interface Towards Provider Edge	Available only for 2-Box VRF. The VRF forwarding interface on the IPsec Aggregator towards the PE device. Click **Select** to select the interface or interface role object, or to create a new object that identifies the interface. **Note** If the IPsec Aggregator (hub) is a Catalyst VPN service module, you must specify a VLAN.
Routing Protocol	Available only for 2-Box VRF. Select the routing protocol to use between the IPsec Aggregator and the PE. The options are BGP, EIGRP, OSPF, RIPv2, or Static route. If the routing protocol for the secured IGP differs from the routing protocol between the IPsec Aggregator and the PE, select the routing protocol for redistributing the routing to the secured IGP.
AS Number	Available only for 2-Box VRF with BGP or EIGRP routing. The number to use to identify the autonomous system (AS) area between the IPsec Aggregator and the PE. The AS number must be between 1 and 65535. If the routing protocol for the secured IGP differs from the routing protocol between the IPsec Aggregator and the PE, enter an AS number that identifies the secured IGP into which the routing will be redistributed from the IPsec Aggregator and the PE. This is relevant only if GRE or DMVPN are applied.
Process Number	Available only for 2-Box VRF with OSPF routing. The routing process ID number to use to configure the routing between the IPsec Aggregator and the PE. The process number must be between 1 and 65535.
OSPF Area ID	Available only for 2-Box VRF with OSPF routing. The ID number of the area in which the packet belongs. You can enter any number from 0 to 4294967295. **Note** All OSPF packets are associated with a single area, so all devices must have the same area ID number.

Table 33-3 IPsec Proposal Editor, Dynamic VTI/VRF Aware IPsec Tab (continued)

Element	Description
Redistribute Static Route	Available only for 2-Box VRF with any routing protocol other than Static route. When selected, enables static routes to be advertised in the routing protocol configured on the IPsec Aggregator towards the PE device. **Note** If this check box is deselected and Enable Reverse Route Injection is enabled (default) for the IPsec proposal, static routes are still advertised in the routing protocol on the IPsec Aggregator.
Next Hop IP Address	Available only for 2-Box VRF with Static routing. The IP address of the provider edge device (or the interface that is connected to the IPSec aggregator).

Configuring High Availability in Remote Access VPNs (IOS)

Use the High Availability page to configure a High Availability (HA) policy on a Cisco IOS router or Cisco Catalyst switch in a remote access VPN.

In Security Manager, High Availability (HA) is supported by the creation of an HA group made up of two or more devices that use Hot Standby Routing Protocol (HSRP) to provide transparent, automatic device failover. By sharing a virtual IP address, the devices in the HA group present the appearance of a single virtual device or default gateway to the remote access users. One device in the HA group is always active and assumes the virtual IP address, while the others are standby devices. The devices in the group watch for hello packets from active and standby devices. If the active device becomes unavailable for any reason, a standby device takes ownership of the virtual IP address and takes over the remote access VPN. This transfer is seamless and transparent to remote access users.

Stateful SwitchOver (SSO) is used to ensure that state information is shared between the HSRP devices in the HA group. If a device fails, the shared state information enables the standby device to maintain IPsec sessions without having to re-establish the tunnel or renegotiate the security associations.

Tips

- When configuring an HA group, you must provide an inside virtual IP that matches the subnet of one of the interfaces on the device, in addition to a VPN virtual IP that matches the subnet of one of the device's interfaces and is configured with an IPsec proposal. See Configuring an IPsec Proposal on a Remote Access VPN Server (IOS, PIX 6.3 Devices), page 33-3.
- A remote access VPN server device on which HA is configured cannot be configured as a hub in a site-to-site VPN topology on which HA is configured, using the same outside interface that was used for the remote access VPN server.

Step 1 Do one of the following:

- (Device view) With an IOS device selected, select **Remote Access VPN > IPSec VPN > High Availability** from the Policy selector.
- (Policy view) Select **Remote Access VPN > IPSec VPN > High Availability** from the Policy Type selector. Select an existing policy or create a new one.

The High Availability page opens.

Step 2 Configure the options explained in the following table.

Table 33-4 High Availability Page, Remote Access VPNs

Element	Description
Inside Virtual IP	The IP address that is shared by the devices in the HA group and that represents the inside interface of the HA group. The virtual IP address must be on the same subnet as the inside interfaces of the devices in the HA group, but must not be identical to the IP address of any of these interfaces. You must provide an inside virtual IP that matches the subnet of one of the interfaces on the device, in addition to a VPN virtual IP that matches the subnet of one of the device's interfaces and is configured with an IPsec proposal. **Note** If there is an existing standby group on the device, make sure that the IP address you provide is different from the virtual IP address already configured on the device.
Inside Mask	The subnet mask for the inside virtual IP address.
VPN Virtual IP	The IP address that is shared by the devices in the HA group and represents the VPN interface of the HA group. This IP address serves as the endpoint of the VPN tunnel. **Note** If there is an existing standby group on the device, make sure that the IP address you provide is different from the virtual IP address already configured on the device.
VPN Mask	The subnet mask for the VPN virtual IP address.
Hello Interval	The duration in seconds (within the range of 1-254) between each hello message sent by a device to the other devices in the group to indicate status and priority. The default is 5 seconds.
Hold Time	The duration in seconds (within the range of 2-255) that a standby device will wait to receive a hello message from the active device before concluding that the device is down. The default is 15 seconds.
Standby Group Number (Inside)	The standby number of the inside device interface that matches the internal virtual IP subnet for the devices in the HA group. The number must be within the range of 0-255. The default is 1.
Standby Group Number (Outside)	The standby number of the outside device interface that matches the external virtual IP subnet for the devices in the HA group. The number must be within the range of 0-255. The default is 2. **Note** The outside standby group number must be different to the inside standby group number.
Failover Server	The IP address or network/host policy object that identifies the inside interface of the remote peer failover servers. Enter the IP address or network/host object name, or click **Select** to select an object or to create a new object.
Enable Stateful Failover	Enables SSO for stateful failover. This option is always selected and you cannot deselect it for remote access VPNs.

Configuring User Group Policies

Use the User Groups (IOS/PIX 6.x) policy to specify user groups for your remote access IPSec VPN server. You can configure user groups on a Cisco IOS router, PIX 6.3 Firewall, or Catalyst 6500 /7600 device.

When you configure a remote access VPN server, you must create user groups to which remote clients will belong. A user group policy specifies the attributes that determine user access to and use of the VPN. User groups simplify system management, enabling you to quickly configure VPN access for large numbers of users.

For example, in a typical remote access VPN, you might allow a finance group to access one part of a private network, a customer support group to access another part, and an MIS group to access other parts. In addition, you might allow specific users within MIS to access systems that other MIS users cannot access. User group policies provide the flexibility to do so securely.

Remote clients must have the same group name as the user group configured on the VPN server so that they can connect to the device; otherwise, a connection cannot be established. When a remote client establishes a connection to the VPN server, the group policies for that user group are pushed to all clients belonging to the same user group. You can configure user groups on the local remote access VPN server and external AAA servers.

Notes

- You can also specify user groups using the Remote Access VPN Configuration Wizard. For more information, see Using the Remote Access VPN Configuration Wizard, page 30-13.
- To specify group policies for an SSL VPN on an IOS device, use the SSL VPN policy as explained in Configuring an SSL VPN Policy (IOS), page 33-14.

Related Topics

- Understanding Remote Access IPSec VPNs, page 30-2

Step 1 Do one of the following:

- (Device view) With an IOS router, Catalyst 6500/7600, or PIX 6.3 device selected, select **Remote Access VPN > IPSec VPN > User Groups (IOS/PIX 6.x)** from the Policy selector.
- (Policy view) Select **Remote Access VPN > IPSec VPN > User Groups (IOS/PIX6.x)** from the Policy Type selector. Select an existing policy or create a new one.

The User Groups page opens.

The page contains two lists: Available User Groups lists all existing User Group policy objects that are configured for remote access IPsec VPNS; Selected User Groups lists all of the User Group policy objects that will be configured on the device.

Step 2 Ensure that the list of selected user groups contains the desired User Group policy objects:

- To create a new User Group policy object, click the Create (+) button beneath the available user groups list to open the Add User Group dialog box. For instructions on creating the object, see Add or Edit User Group Dialog Box, page 34-73.

  After you create the group, it is added to the available list, and you must add it to the selected list if you want to use it.
- To add a User Group to the selected list, select it in the available list and click >>.
- To remove a User Group, select it in the selected list and click <<. If the group is already configured on the device, it will be removed during the next deployment.

- You can edit the properties of a User Group object by selecting it in either list and clicking the **Edit** button.

Configuring an SSL VPN Policy (IOS)

Use the SSL VPN policy to configure the SSL VPN connection policies for an IOS router. From this page, you can create, edit, or delete SSL VPN policies.

Related Topics

- Understanding Remote Access SSL VPNs, page 30-2
- Creating SSL VPNs Using the Remote Access VPN Configuration Wizard (IOS Devices), page 30-32
- Filtering Tables, page 1-48

Step 1 Do one of the following:

- (Device view) With an IOS device selected, select **Remote Access VPN > SSL VPN** from the Policy selector.
- (Policy view) Select **Remote Access VPN > SSL VPN > SSL VPN Policy (IOS)** from the Policy Type selector. Select an existing policy or create a new one.

The SSL VPN page appears.

The table lists all of the contexts that define the virtual configurations of the SSL VPN. Each context has a gateway, domain or virtual hostname, and user group policies. The status of the context is also shown, either In Service or Out of Service.

Step 2 Do either of the following:

- To add a context, click the **Add Row** button to open the SSL VPN Context Editor Dialog Box (IOS), page 33-15.
- To edit a context, select it and click the **Edit Row** button.

Note To delete a context, select it and click the **Delete Row** button.

Step 3 Configure at least the following general settings for the policy. For information on other fields, see General Tab, page 33-16.

- **Name, Domain**—For new policies, the name of the context that defines the virtual configuration of the SSL VPN. To simplify the management of multiple context configurations, make the context name the same as the domain or virtual hostname.
- **Gateway**—The SSL VPN gateway policy object that identifies the gateway device to which users will connect, including interface and port configuration. Click **Select** to select the object from a list or to create a new object.

  When you select the object, the Portal Page URL field shows the URL to which users connect.
- **Authentication Server Group**—A prioritized list of AAA server group objects that identify the AAA servers to use for authenticating users.

- **User Groups**— The user groups that will be used in your SSL VPN policy. User groups define the resources available to users when connecting to an SSL VPN gateway.

  To add a user group, click **Add Row** to open a list of existing user group policy objects from which you can select the group. If the desired group does not already exist, click the **Create** button below the available groups list and create it. For more information about user group objects, see Add or Edit User Group Dialog Box, page 34-73.

Step 4 Click the **Portal Page** tab and customize the design of the login page. You can customize the title, the logo graphic, the message that appears above the login prompt, and the background and text colors.

If you want to select a different graphic, you must first copy the graphic onto the Security Manager server. You cannot select it from your workstation's hard drive.

Step 5 Click the **Secure Desktop** tab to configure Cisco Secure Desktop (CSD) software. CSD policies define entry requirements for client systems and provide a single, secure location for session activity and removal on the client system, ensuring that sensitive data is shared only for the duration of an SSL VPN session.

If you want to use CSD, select **Enable Cisco Secure Desktop** and click **Select** to select a Secure Desktop Configuration policy object, which defines the rules you want to use to control VPN access and host scanning. You can create a new object from the selection list. For information about configuring these objects, see Creating Cisco Secure Desktop Configuration Objects, page 33-18.

Note You must install and activate the Secure Desktop Client software on a device for your configuration to work.

Step 6 Click the **Advanced** tab to configure a maximum number of simultaneous users for the context or if you are using VRF, the name of the VRF instance that is associated with the SSL VPN context.

Step 7 Click **OK** to save your changes.

SSL VPN Context Editor Dialog Box (IOS)

Use this dialog box to create or modify a context that defines the virtual configuration of an SSL VPN. For more information, see Configuring an SSL VPN Policy (IOS), page 33-14.

Navigation Path

Open the SSL VPN (IOS) policy, then click **Add Row (+)**, or select a context in the table and click **Edit Row**. For information on opening the SSL VPN policy, see Configuring an SSL VPN Policy (IOS), page 33-14.

Field Reference

Table 33-5 SSL VPN Context Editor Dialog Box

Element	Description
General tab	Defines the general settings required for an SSL VPN policy. General settings include specifying the gateway, domain, AAA servers for accounting and authentication, and user groups. For a description of the fields on this tab, see General Tab, page 33-16.

Table 33-5 SSL VPN Context Editor Dialog Box (continued)

Element	Description
Portal Page tab	Defines the design of the login page for the SSL VPN policy. The display box at the bottom of the tab changes to show you how your selections will look. You can configure: • Title—The text displayed at the top of the page. Control the color using the Primary settings in the Title Color and Text Color fields. • Logo—The graphic displayed next to the title. Select None, Default, or Custom. To configure a custom graphic, you must copy the desired graphic to the Security Manager server, then click **Browse** to select the file. Supported graphic types are GIF, JPG, and PNG, with a maximum size of 100 KB. • Login Message—The text displayed immediately above the login prompt. Control the color using the Secondary settings in the Title Color and Text Color fields.
Secure Desktop tab	Configures the Cisco Secure Desktop (CSD) software on the router. CSD policies define entry requirements for client systems and provide a single, secure location for session activity and removal on the client system, ensuring that sensitive data is shared only for the duration of an SSL VPN session. **Note** You must install and activate the Secure Desktop Client software on a device for your configuration to work. If you want to use CSD, select **Enable Cisco Secure Desktop** and click **Select** to select a Secure Desktop Configuration policy object, which defines the rules you want to use to control VPN access and host scanning. You can create a new object from the selection list. For information about configuring these objects, see Creating Cisco Secure Desktop Configuration Objects, page 33-18.
Advanced tab	Configures these additional settings: • Maximum Number of Users—The maximum number of SSL VPN user sessions allowed at one time, from 1-1000. • VRF Name—If Virtual Routing Forwarding (VRF) is configured on the device, the name of the VRF instance that is associated with the SSL VPN context. For information about VRF, see Understanding VRF-Aware IPsec, page 25-14.

General Tab

Use the General tab of the SSL VPN Context Editor dialog box to define or edit the general settings required for an SSL VPN policy. General settings include specifying the gateway, domain, AAA servers for accounting and authentication, and user groups.

Navigation Path

Open the SSL VPN Context Editor Dialog Box (IOS), page 33-15, then click the **General** tab.

Related Topics

- Configuring an SSL VPN Policy (IOS), page 33-14
- Add or Edit SSL VPN Gateway Dialog Box, page 34-64
- Understanding AAA Server and Server Group Objects, page 6-27

Field Reference

Table 33-6 SSL VPN Context Editor General Tab (IOS)

Element	Description
Enable SSL VPN	Whether to activate the SSL VPN connection, putting it "In Service".
Name	The name of the context that defines the virtual configuration of the SSL VPN. **Note** To simplify the management of multiple context configurations, make the context name the same as the domain or virtual hostname.
Gateway	The name of the SSL VPN gateway policy object that defines the characteristics of the gateway to which users connect when entering the VPN. A gateway object provides the interface and port configuration for an SSL VPN connection. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Domain	The domain or virtual hostname of the SSL VPN connection.
Portal Page URL	The URL for the SSL VPN, which is filled in when you select a gateway object. Users connect to this URL to enter the VPN.
Authentication Server Group	The authentication server groups. The list is in prioritized order. Authentication is attempted using the first group and proceeds through the list until the user is successfully authenticated or denied. Use the LOCAL group if the users are defined on the gateway itself. Enter the names of the AAA server groups; separate multiple entries with commas. You can click **Select** to select the groups or to create new ones.
Authentication Domain	A list or method for SSL VPN remote user authentication. If you do not specify a list or method, the gateway uses global AAA parameters for remote-user authentication.
Accounting Server Group	The accounting server group. Enter the name of the AAA server group policy object, or click **Select** to select it from a list or to create a new object.

Table 33-6 SSL VPN Context Editor General Tab (IOS) (continued)

Element	Description
User Groups	The user groups that will be used in your SSL VPN policy. User groups define the resources available to users when connecting to an SSL VPN gateway. The table shows whether full client, CIFS file access, and thin client is enabled for the group. • To add a user group, click **Add Row** to open a list of existing user group policy objects from which you can select the group. If the desired group does not already exist, click the **Create** button below the available groups list and create it. For more information about user group objects, see Add or Edit User Group Dialog Box, page 34-73. • To edit a user group, select it and click the **Edit Row** button. • To delete a user group, select it and click the **Delete Row** button. This deletes the group only from the policy, it does not delete the user group policy object.

Creating Cisco Secure Desktop Configuration Objects

Cisco Secure Desktop (CSD) Configuration objects define the settings you want to use if you enable Secure Desktop in an SSL VPN policy for an IOS device (see Configuring an SSL VPN Policy (IOS), page 33-14). For ASA devices, the feature is set up as part of the Dynamic Access Policy (see Understanding Dynamic Access Policies, page 32-1 and Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9).

Cisco Secure Desktop (CSD) provides a reliable means of eliminating all traces of sensitive data by providing a single, secure location for session activity and removal on the client system. CSD provides a session-based interface where sensitive data is shared only for the duration of an SSL VPN session. All session information is encrypted, and all traces of the session data are removed from the remote client when the session is terminated, even if the connection terminates abruptly.

About Windows Locations

Windows locations let you determine how clients connect to your virtual private network, and protect it accordingly. For example, clients connecting from within a workplace LAN on a 10.x.x.x network behind a NAT device are an unlikely risk for exposing confidential information. For these clients, you might set up a CSD Windows Location named Work that is specified by IP addresses on the 10.x.x.x network, and disable both the Cache Cleaner and the Secure Desktop function for this location.

In contrast, users' home PCs might be considered more at risk to viruses due to their mixed use. For these clients, you might set up a location named Home that is specified by a corporate-supplied certificate that employees install on their home PCs. This location would require the presence of antivirus software and specific, supported operating systems to grant full access to the network.

Alternatively, for untrusted locations such as Internet cafes, you might set up a location named "Insecure" that has no matching criteria (thus making it the default for clients that do not match other locations). This location would require full Secure Desktop functions, and include a short timeout period to prevent access by unauthorized users. If you create a location and do not specify criteria, make sure it is the last entry in the Locations list.

Related Topics

- Cisco Secure Desktop on IOS Configuration Example Using SDM, http://www.cisco.com/en/US/products/ps6496/products_configuration_example09186a008072aa7b.shtml
- Setting Up CSD for Microsoft Windows Clients, http://www.cisco.com/en/US/docs/security/csd/csd311/csd_for_vpn3k_cat6k/configuration/guide/CSDwin.html
- Creating Policy Objects, page 6-9

Step 1 Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4).

Step 2 Select **Cisco Secure Desktop Configuration** from the Object Type selector.

Step 3 Right-click in the work area and select **New Object** to open the Add or Edit Secure Desktop Configuration Dialog Box, page 34-35.

Step 4 Enter a name for the object and optionally a description of the object.

Step 5 Select **Windows Location Settings** to create locations (such as Work, Home, or Insecure), and define the location-based settings (also called adaptive policies) for CSD.

a. For each location you want to configure, enter its name in the **Location to Add** field and click **Add** to move it to the Locations field. You can reorder the locations using the Move Up and Move Down buttons. When users connect, these locations are evaluated in order and the first one that matches is used to define the policies for the user.

When you add a location, a folder for the location is added to the table of contents. The folder and its subfolders define the policies for the location.

b. If you want all the open browser windows to close after the Secure Desktop installation, make sure to select the corresponding check box.

c. Select the required check boxes to configure a VPN Feature policy that enables web browsing, file access, port forwarding, and full tunneling, if installation or location matching fails.

Step 6 Select the folders and subfolders for the Windows locations you added and configure their settings. For detailed information about these settings, see *Setting Up CSD for Microsoft Windows Clients* at http://www.cisco.com/en/US/docs/security/csd/csd311/csd_for_vpn3k_cat6k/configuration/guide/CSDwin.html.

Step 7 Select **Windows CE** to configure a VPN feature policy to enable or restrict web browsing and remote server file access for remote clients running Microsoft Windows CE.

Step 8 Select **Mac and Linux Cache Cleaner** to configure the Cache Cleaner and a VPN Feature Policy for these clients, such as enabling or restricting web browsing, remote server file access, and port forwarding.

Step 9 (Optional) Under Category, select a category to help you identify this object in the Objects table. See Using Category Objects, page 6-13.

Step 10 Click **OK** to save the object.

Configuring Policy Objects for Remote Access VPNs

There are several policy objects that you use primarily or exclusively with remote access VPNs. Some of these objects, the ASA Group Policies and User Group objects, are also used with Easy VPN site-to-site topologies. This reference explains the configuration of these policy objects.

This chapter contains the following topics:

- ASA Group Policies Dialog Box, page 34-1
- Add or Edit Secure Desktop Configuration Dialog Box, page 34-35
- Add and Edit File Object Dialog Boxes, page 34-37
- Add or Edit Port Forwarding List Dialog Boxes, page 34-40
- Add or Edit Single Sign On Server Dialog Boxes, page 34-42
- Add or Edit Bookmarks Dialog Boxes, page 34-44
- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51
- Add or Edit SSL VPN Gateway Dialog Box, page 34-64
- Add and Edit Smart Tunnel List Dialog Boxes, page 34-66
- Add and Edit Smart Tunnel Network Lists Dialog Boxes, page 34-69
- Add and Edit Smart Tunnel Auto Signon List Dialog Boxes, page 34-71
- Add or Edit User Group Dialog Box, page 34-73
- Add or Edit WINS Server List Dialog Box, page 34-89

ASA Group Policies Dialog Box

Use the Add or Edit ASA Group Policies dialog box to create, copy, and edit an ASA user group policies object.

ASA group policies are configured on ASA security appliances in Easy VPN topologies, remote access IPSec VPNs, and remote access SSL VPNs. When you configure an Easy VPN or remote access VPN, you must create group policies to which remote clients will belong. A group policy is a set of user-oriented attribute/value pairs for VPN connections that are stored either internally (locally) on the device or externally on a AAA server. The tunnel group uses a group policy that sets terms for user connections after the tunnel is established. Group policies let you apply whole sets of attributes to a user or a group of users rather than having to specify each attribute individually for each user.

Note You must select the technology for which you are creating the object. Depending on the selected technology, the appropriate settings are available for configuration. If you select the IKEv1 or IKEv2 options, the IKE Proposal and IPSec Proposal policies must also be configured to support the selected IKE version.

From version 4.18, Cisco Security Manager has introduced the option to override group policies. In the ASA Group Policy page, you can enable device overrides and edit device overrides from right-click menu. When override is enabled,

Navigation Path

Select **ASA Group Policies** in the Policy Object Manager, page 6-4. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Tip You can also create objects while configuring policies that use this type of object, including Connection Profile policies for remote access and Easy VPN, or the Group Policies policy for remote access VPNs.

Related Topics

- Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7
- Creating Group Policies (ASA, PIX 7.0+), page 31-28

Field Reference

Table 34-1 Add or Edit ASA Group Policies Dialog Box, including Technology Settings

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Settings Pane

The body of the dialog box is a pane with a table of contents on the left and settings related to the item selected in the table of contents on the right.

You must first configure technology settings, then you can select items from the table of contents on the left and configure the options you require. Your selections on the Technology page control which options are available on these pages and in the table of contents.

The top folders in the table of contents represent the VPN technologies or other settings that you can configure, and are explained next.

Table 34-1 Add or Edit ASA Group Policies Dialog Box, including Technology Settings

Element	Description
Technology settings	These settings control what you can define in the group policy: • **Group Policy Type**—Whether you are storing the group policy on the ASA device itself (**Internal**) or on a AAA server (**External**). You cannot change this option when editing an object. If you select **External**, the only attributes you can configure are the name of the AAA server group object that identifies the AAA server and its password. • **Technology**—The types of VPN for which this object defines group policies. Select all that apply: – Easy VPN/IPSec IKEv1—For Easy VPN topologies or remote access IPsec VPNs that allow IKEv1 negotiations. – Easy VPN/IPSec IKEv2—For remote access IPsec VPNs that allow IKEv2 negotiations. IKEv2 is not supported in Easy VPN topologies. – SSL Clientless—For remote access SSL VPNs of all types, not just clientless. **Note** To enable web-based VPN (webvpn) option in group-policy attribute, you must enable either "ssl-client" or "ssl-clientless" tunneling protocol. In other words, upon device discovery in Security Manager, if the group-policy attribute "vpn-tunnel-protocol" does not have either "ssl-client" or "ssl-clientless" in the configuration, during the next deployment of the device, Security Manager would remove the "webvpn" option under group-policy attributes. • **External Server Group**—If you are storing the group policy attributes on an external AAA server, specify the AAA server group that will be used for authentication. Click **Select** to select the object from a list or to create a new object. After you select an external server group, the **Password** and **Confirm** fields become active. Enter the alphanumeric password to use for authenticating with the server in both fields. The password can be a maximum of 128 characters; spaces are not allowed.
DNS/WINS	The DNS and WINS servers and the domain name that should be pushed to clients associated with the group. See ASA Group Policies DNS/WINS Settings, page 34-30.
Split Tunneling	Settings to allow a remote client to conditionally direct encrypted packets through a secure tunnel to the central site and simultaneously allow clear text tunnels to the Internet through a network interface. See ASA Group Policies Split Tunneling Settings, page 34-31.

Table 34-1 Add or Edit ASA Group Policies Dialog Box, including Technology Settings

Element	Description
Easy VPN/IPSec VPN	Settings for Easy VPN and remote access IPSec VPNs: • Client Configuration—The Cisco client parameters for the group. See ASA Group Policies Client Configuration Settings, page 34-6. • Client Firewall Attributes—The firewall settings for VPN clients for the group. See ASA Group Policies Client Firewall Attributes, page 34-7. • Hardware Client Attributes—The VPN 3002 Hardware Client settings for the group. See ASA Group Policies Hardware Client Attributes, page 34-9. • IPSec—The tunneling protocols, filters, connection settings, and servers for the group. See ASA Group Policies IPSec Settings, page 34-10.
SSL VPN	Settings for SSL VPN: • Clientless—Settings for the clientless mode of access to the corporate network in an SSL VPN. See ASA Group Policies SSL VPN Clientless Settings, page 34-12. • Full Client—Settings for the full client mode of access to the corporate network in an SSL VPN. See ASA Group Policies SSL VPN Full Client Settings, page 34-19. • Settings—The general settings that are required for clientless/port forwarding in an SSL VPN. See ASA Group Policies SSL VPN Settings, page 34-25.
Connection Settings	The connection settings for the group, such as the session and idle timeouts, including the banner text. See ASA Group Policies Connection Settings, page 34-33.
General Settings	• Override Group Policy—Beginning with version 4.18, Cisco Security Manager allows group policy overrides. See Override ASA Group Policy, page 34-4.

Override ASA Group Policy

In Cisco Security Manager, group policies are created for the devices and maintained at the Cisco Security Manager level. When there is an upgrade, on rediscovery, Cisco Security Manager recreates these policies as new (with a suffix to the policy name). To overcome this duplication, from version 4.18, an Allow Value Override per device check box is used to set the group policy override on the specific device(s). For more information, see Managing Object Overrides, page 6-17.

You can edit the device-level overrides for the group policies. See Policy Object Overrides Window, page 6-20.

Supported CLIs in Remote Access VPN Multi-Context Mode - Group Policy

The following CLIs are supported for Group Policy in ASA version 9.5(2) for remote access VPN in multiple context mode. These CLIs are supported in Admin and User Context.

Note For the configurations that are not supported, Security Manager displays a warning message that you can ignore. No delta will be generated.

- Address-pools
- Banner
- Client-bypass-protocol
- Default-domain
- Dhcp-network-scope
- Dns-server
- Exit
- Gateway-fqdn
- Gateway-fqdn
- Ipv6-address-pools
- Ipv6-address-pools
- Msie-proxy
- No
- Security-group-tag
- Smartcard-removal-disconnect
- Periodic-authentication
- Split-dns
- Split-tunnel-all-dns
- Split-tunnel-network-list
- Split-tunnel-policy
- Vpn-access-hours
- Vpn-filter (already supported in multi-mode for S2S)
- Vpn-simultaneous-logins
- Vpn-idle-timeout (already supported in multi-mode for S2S)
- Vpn-session-timeout (already supported in multi-mode for S2S)
- Vpn-tunnel-protocol ssl-client
- Wins-server
- Webvpn
  - Anyconnect-custom
  - anyconnect Dpd-interval
  - anyconnect dtls

- anyconnect firewall-rule
- anyconnect keep-installer
- anyconnect modules
- anyconnect Mtu
- anyconnect routing-filtering-ignore
- anyconnect Ssl
- exit
- homepage value | none
- no

ASA Group Policies Client Configuration Settings

Use the Client Configuration settings page to configure the Cisco client parameters for the ASA group policy for Easy VPN or remote access VPN.

Client Configuration is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Easy VPN/IPSec VPN > Client Configuration** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-2 ASA Group Policies Client Configuration Settings

Element	Description
Store Password on Client System	Whether to allow users to store a password on their local systems. Enable this feature only if you are certain that the local systems will be in secure sites.
Enable IPsec over UDP UDP Port	Whether to allow a Cisco VPN client or hardware client to connect using UDP to a security appliance that is running NAT. If you select this option, specify the UDP port number within the range of 4001-49151. In IPsec negotiations, the security appliance listens on the configured port and forwards UDP traffic for that port even if other filter rules drop UDP traffic. **Note** The Cisco VPN client must also be configured to use IPsec over UDP, which is configured by default on certain devices.

Table 34-2 ASA Group Policies Client Configuration Settings (continued)

Element	Description
IPsec Backup Servers Servers List	Specify the backup server configuration: • **Keep Client Configuration**—The security appliance sends no backup server information to the client. The client uses its own backup server list, if configured. This is the default. • **Clear Client Configuration**—The client uses no backup servers. The security appliance pushes a null server list. • **Use Specified Backup Servers**—Use the backup servers you specify in the servers list. Enter the IP addresses of the servers, or the name of a network/host object. Click **Select** to select the object from a list or to create a new object. You can configure backup servers either on the client or on the primary security appliance. If you configure backup servers on the security appliance, it pushes the backup server policy to the clients in the group, replacing the backup server list on the client if one is configured.

ASA Group Policies Client Firewall Attributes

Use the Client Firewall Attributes settings to configure the firewall settings for VPN clients for the ASA group policy for Easy VPN or remote access IPSec VPN. Only VPN clients running Microsoft Windows can use these firewall settings.

Client Firewall Attributes are not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Easy VPN/IPSec VPN > Client Firewall Attributes** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-3 ASA Group Policies Client Firewall Attributes

Element	Description
Firewall Mode	The firewall requirements for client systems for the group: • **No Firewall**—Do not use a firewall. You cannot configure any other options on the page. • **Firewall Required**—All users in this group must use the designated firewall. The security appliance drops any session that attempts to connect without the designated firewall installed and running. In this case, the security appliance notifies the VPN client that its firewall configuration does not match. **Note** Make sure the group does not include any clients other than Windows VPN Clients. Any other clients in the group (including VPN 3002 Hardware Clients) are unable to connect if you require a client firewall. • **Firewall Optional**—Users can use a firewall but it is not required. This option allows all users in the group to connect. Those who have a firewall can use it; users that connect without a firewall receive a warning message. This setting is useful if you are creating a group in which some users have firewalls and others do not. For example, you might have clients with systems that do not run Microsoft windows, or your clients have not all had the opportunity to install firewall software.
Firewall Type	The type of firewall that you are making required or optional. The list shows all of the supported firewall software, which includes software from Cisco, Network ICE, Sygate, and Zone Labs. • If you select Custom Firewall, you must fill in the fields in the Custom Firewall group. You also need to configure the policy source; select options only if they are supported by the vendor. • Some firewall types require you to specify the source of the policy implemented by the firewall.
Policy Source	Some types of firewall allow you to configure where the client firewall should obtain its policies: • Get Policy From Remote Firewall—The policy is configured in the client firewall application. This is how most client firewalls work. • Use Specified Policy—The policy you specify should be pushed to the client firewall application, which should use your policy. You must enter the name of an extended access control list policy object or Unified ACL, or click **Select** to select one from a list or to create a new one, in both in the **Inbound Traffic Policy** and **Outbound Traffic Policy** fields. Unified ACLs are supported from ASA version 9.0.

Table 34-3 ASA Group Policies Client Firewall Attributes (continued)

Element	Description
Custom Firewall	The attributes that define the required or optional firewall if you select custom firewall as the firewall type: • Vendor ID—The number that identifies the vendor of the custom firewall. Values are 1-255. • Product ID—The number that identifies the product or model of the custom firewall. Values are 1-32 or 255. Multiple ranges are allowed, for example, 4-12, 24-32. Use 255 for all supported products. • Description—An optional description of the custom firewall, for example, the name of the vendor and product.

ASA Group Policies Hardware Client Attributes

Use the Hardware Client Attributes settings to configure the VPN 3002 Hardware Client settings for the ASA group policy in an Easy VPN or remote access IPSec VPN.

Hardware Client Attributes are not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Easy VPN/IPSec VPN > Hardware Client Attributes** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-4 ASA Group Policies Hardware Client Attributes

Element	Description
Require Interactive Client Authentication	Whether to enable secure unit authentication, which provides additional security by requiring VPN hardware clients to authenticate with a username and password each time that the client initiates a tunnel. The hardware client does not have a saved username and password. **Note** Secure unit authentication requires that you have an authentication server group configured for the tunnel group the hardware clients use. If you require secure unit authentication on the primary security appliance, be sure to configure it on any backup servers as well.
Require Individual User Authentication	Whether to require that individual users behind a hardware client authenticate to gain access to the network across the tunnel. Individual users authenticate according to the order of authentication servers that you configure. If you do not select this option, the security appliance allows inheritance of a value for user authentication from another group policy.

Table 34-4 ASA Group Policies Hardware Client Attributes (continued)

Element	Description
Enable Cisco IP Phone Bypass	Whether to allow IP phones behind hardware clients to connect without undergoing a user authentication processes. Secure unit authentication remains in effect for other users.
Enable LEAP Bypass	Whether to enable Lightweight Extensible Authentication Protocol (LEAP) packets from wireless devices behind a VPN hardware client to travel across a VPN tunnel prior to user authentication. This action lets workstations using Cisco wireless access point devices establish LEAP authentication and then authenticate again per user authentication. **Note** LEAP is an 802.1X wireless authentication method that implements mutual authentication between a wireless client on one side of a connection and a RADIUS server on the other side. The credentials used for authentication, including a password, are always encrypted before they are transmitted over the wireless medium.
Allow Network Extension Mode	Whether to enable network extension mode for hardware clients. Network extension mode lets hardware clients present a single, routable network to the remote private network over the VPN tunnel. IPsec encapsulates all traffic from the private network behind the hardware client to networks behind the security appliance. PAT does not apply. Devices behind the security appliance have direct access to devices on the private network behind the hardware client over the tunnel, and only over the tunnel, and vice versa. The hardware client must initiate the tunnel, but after the tunnel is up, either side can initiate data exchange.
Idle Timeout Mode	How to handle periods of inactivity from individual clients: • Specified Timeout—If there is no communication activity by a user behind a hardware client for the number of minutes you specify, the security appliance terminates the client's access. Values are 1-35791394 minutes. • Unlimited Timeout—User sessions are not terminated due to inactivity.

ASA Group Policies IPSec Settings

Use the IPsec settings to specify tunneling protocols, filters, connection settings, and servers for the ASA group policy for Easy VPN or remote access IPSec VPN. This creates security associations that govern authentication, encryption, encapsulation, and key management.

IPSec is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Easy VPN/IPSec VPN > IPsec** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-5 ASA Group Policies IPSec Settings

Element	Description
Enable Re-Authentication on IKE Re-Key	Whether the security appliance should prompt the user to enter a username and password during initial Phase 1 IKE negotiation and also prompt for user authentication whenever an IKE rekey occurs, providing additional security. Reauthentication fails if no user is at the other end of the connection.
Enable IPsec Compression	Whether to enable data compression, which speeds up transmission rates for remote dial-in users connecting with modems. **Caution** Data compression increases the memory requirement and CPU usage for each user session and consequently decreases the overall throughput of the security appliance. For this reason, it is recommended that you enable data compression only for remote users connecting with a modem. Design a group policy specific to modem users and enable compression only for them.
Enable Perfect Forward Secrecy (PFS)	Whether to enable the use of Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. In IPsec negotiations, PFS ensures that each new cryptographic key is unrelated to any previous key.
Tunnel Group Lock	Tunnel group lock restricts users by checking if the group configured in the VPN client is the same as the tunnel group to which the user is assigned. If it is not, the security appliance prevents the user from connecting. If you do not specify a tunnel name, the security appliance authenticates users without regard to the assigned group. Group locking is disabled by default.
Client Access Rules table	The access rules for clients. These rules control which types of clients are denied access, if any. You can have up to 25 rules, and combined they are limited to 255 characters. **Tip** If you define any rule, an implicit deny all rule is added. Thus, if a client matches no permit rule, the client is denied access. If you create rules, ensure that you have permit rules for all allowed clients. You can use * as a wildcard to match partial strings. The rule with the lowest integer has the highest priority. Therefore, the rule with the lowest integer that matches a client type or version is the rule that applies. If a lower priority rule contradicts, the security appliance ignores it. • To add a rule, click the **Add Row** button to open the Add or Edit Client Access Rules Dialog Box, page 34-12. • To edit a rule, select it and click the **Edit Row** button. • To delete a rule, select it and click the **Delete** button.

Add or Edit Client Access Rules Dialog Box

Use the Client Access Rules dialog box to create or edit the priority, action, VPN client type and VPN client version for a client access rule.

Navigation Path

From ASA Group Policies IPSec Settings, page 34-10, click the **Add Row** button beneath the Client Access Rules table, or select a rule and click the **Edit Row** button.

Field Reference

Table 34-6 Add or Edit Client Access Rules Dialog Box

Element	Description
Priority	The relative priority of the rule. The rule with the lowest integer has the highest priority. Therefore, the rule with the lowest integer that matches a client type or version is the rule that applies. If a lower priority rule contradicts, the security appliance ignores it. Values are 1-65535.
Action	Whether this rule permits or denies traffic access to the client.
VPN Client Type VPN Client Version	The type or version of VPN client to which this rule applies. Spaces are allowed. You can use * as a wildcard to match zero or more characters. You can use n/a for clients that do not send their type or version. The strings you enter in these fields must match the strings displayed using the **show vpn-sessiondb remote** command on the ASA device. Following are some examples, where priority, permit/deny, type, and version are shown in order: • **3 Deny * version 3.*** is a priority 3 rule that denies all client types with software version 3.x. • **5 Permit VPN3002 *** is a priority 5 rule that allows VPN3002 clients of all software versions. • **255 Permit * *** is a priority 255 rule that allows all types and versions of clients. This is useful if you are only trying to deny specific types of clients without wanting to create permit rules for all the other types.

ASA Group Policies SSL VPN Clientless Settings

Use the Clientless settings to configure the clientless mode of access to the corporate network in a remote access SSL VPN for the ASA group policy object.

When a user connects to the SSL VPN in clientless mode, the user logs into the SSL VPN portal page. From the portal page, the user can access all available HTTP sites, access web e-mail, and browse Common Internet File System (CIFS) file servers, depending on how you configure the portal.

Clientless is not supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **SSL VPN > Clientless** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-7 ASA Group Policies SSL VPN Clientless Settings

Element	Description
Portal Page Websites	The name of the SSL VPN bookmarks policy object that includes the website URLs to display on the portal page. These websites help users access desired resources. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Allow Users to Enter Websites	Whether to allow the remote user to enter website URLs directly into the browser. If you do not select this option, the user can access only those URLs included on the portal.
Enable File Server Browsing	Whether to allow the remote user to browse for file shares on the CIFS file servers.
Enable File Server Entry	Whether to allow the remote user to locate file shares on the CIFS file servers by entering the names of the file shares.
Enable Hidden Shares	Whether to make hidden CIFS shares visible, and thus accessible, to users.
HTTP Proxy	The type of access you want to allow to the external HTTP proxy server to which the security appliance forwards HTTP connections. You can enable access, disable access, or select Auto Start, which starts the proxy automatically upon user login.
Filter ACL	The name of the web type access control list policy object to use to restrict user access to the SSL VPN. Enter the name of the object or click **Select** to select it from a list or to create a new object. Beginning with version 4.10, you can enter IPv6 values for the web type ACL.
Enable ActiveX Relay	Whether to enable ActiveX relay, which allows users to start ActiveX programs from the portal page. This allows users to start Microsoft Office applications from the web browser and upload and download Office documents.
UNIX Authentication Group ID	The UNIX authentication group ID.
UNIX Authentication User ID	The UNIX authentication user ID.
Smart Tunnel	The name of the smart tunnel list policy object assigned to this group. Click **Select** to select it from a list or to create a new object. A smart tunnel is a connection between a Winsock 2, TCP-based application and a private site. The connection uses a clientless (browser-based) SSL VPN session with the security appliance as the pathway, and the security appliance as a proxy server. Thus, smart tunnels do not require users to have administrator privileges. For more information, see Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85.

Table 34-7 ASA Group Policies SSL VPN Clientless Settings (continued)

Element	Description
Auto Start Smart Tunnel	Whether to start smart tunnel access automatically upon user login. If you do not select this option, the user must start the tunnel manually through the Application Access tools on the portal page. Auto sign-on supports only applications that use HTTP and HTTPS using the Microsoft WININET library on a Microsoft Windows operating system. For example, Microsoft Internet Explorer uses the WININET dynamic linked library to communicate with web servers.
Smart Tunnel Network List	Choose from the following options to select the list of hosts or network for which you want to use the smart tunnel. To enable the selection, you must first create the smart tunnel network list entries. For more information, see Add and Edit Smart Tunnel Network Lists Dialog Boxes, page 34-69. Note that this feature is supported on devices that are running ASA software version 8.3(1) and higher. • None—If you select this option, the group policy inherits the values from the default Group Policy. This option is enabled by default. • Tunnel All—Select this option if you want to use the smart tunnel for all network traffic. • Include—Select this option if you want to use the smart tunnel for specific networks. Then click **Select** to open the Smart Tunnel Network List Selector dialog box. You can select from the available entries or add entries. To add smart tunnel network list entries, see Add and Edit Smart Tunnel Network Lists Dialog Boxes, page 34-69. • Exclude—Select this option if you do not want to use the smart tunnel for specific networks. Then click **Select** to open the Smart Tunnel Network List Selector dialog box. You can select from the available entries or add entries. To add smart tunnel network list entries, see Add and Edit Smart Tunnel Network Lists Dialog Boxes, page 34-69.
Smart Tunnel Auto Signon Server List	The name of the smart tunnel auto sign-on list policy object assigned to this group. Click **Select** to select it from a list or to create a new object.
Domain Name (Optional)	The Windows domain to add to the username during auto sign-on, if the universal naming convention (domain\username) is required for authentication. For example, enter CISCO to specify CISCO\qa_team when authenticating for the username qa_team. You must also check the Use Domain option when configuring associated entries in the auto sign-on server list.
Port Forwarding List	The name of the port forwarding list policy object assigned to this group. Port forwarding lists contain the set of applications that users of clientless SSL VPN sessions can access over forwarded TCP ports. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Auto Start Port Forwarding	Whether to start port forwarding automatically upon user login.

Table 34-7 ASA Group Policies SSL VPN Clientless Settings (continued)

Element	Description
Port Forwarding Applet Name	The application name or short description to display on the Port Forwarding Java applet screen on the portal, up to 64 characters. This is the name of the applet users will download to act as a TCP proxy on the client machine for the services configured on the SSL VPN gateway.
VDI Servers List table	The Citrix XenApp or XenDesktop servers that comprise the Virtual Desktop Infrastructure. • To add a VDI server, click the **Add Row** button to open the Add or Edit VDI Server Dialog Box, page 34-15. • To edit a rule, select it and click the **Edit Row** button. • To delete a rule, select it and click the **Delete** button.

Add or Edit VDI Server Dialog Box

Use the VDI Server dialog box to create or edit a Citrix XenApp or XenDesktop Server entry.

In a Virtual Desktop Infrastructure (VDI) model, administrators publish enterprise applications or desktops pre-loaded with enterprise applications, and end users remotely access these applications. These virtualized resources appear just as any other resources, such as email, so that users do not need to go through a Citrix Access Gateway to access them. Users log onto the ASA using Citrix Receiver mobile client, and the ASA connects to a pre-defined Citrix XenApp or XenDesktop Server. The administrator must configure the Citrix server's address and logon credentials under Group Policy so that when users connect to their Citrix Virtualized resource, they enter the ASA's SSL VPN IP address and credentials instead of pointing to the Citrix Server's address and credentials. When the ASA has verified the credentials, the receiver client starts to retrieve entitled applications through the ASA.

Supported Mobile Devices

- iPad—Citrix Receiver version 4.x or later
- iPhone/iTouch—Citrix Receiver version 4.x or later
- Android 2.x/3.x/4.0/4.1 phone—Citrix Receiver version 2.x or later
- Android 4.0 phone—Citrix Receiver version 2.x or later

Navigation Path

From ASA Group Policies SSL VPN Clientless Settings, page 34-12, click the **Add Row** button beneath the VDI Servers List table, or select a rule and click the **Edit Row** button.

Field Reference

Table 34-8 Add or Edit VDI Server Dialog Box

Element	Description
Hostname/IP Address (IPv4/IPv6)	Address of the XenApp or XenDesktop server. This value can be a clientless macro. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running the version 9.0 or later. For invalid IPv6 addresses, Security Manager throws up an error.

Table 34-8 Add or Edit VDI Server Dialog Box (continued)

Element	Description
Port Number (Optional)	Port number for connecting to the Citrix server. This value can be a clientless macro.
Domain	Domain for logging into the virtualization infrastructure server. This value can be a clientless macro.
Secure HTTP	Check the checkbox if you want the server to connect using SSL.

Table 34-8 Add or Edit VDI Server Dialog Box (continued)

Element	Description
Username	Username for logging into the virtualization infrastructure server. This value can be a clientless macro. The macros available for username are: • CSCO_WEBVPN_USERNAME—SSL VPN user login ID. • CSCO_WEBVPN_CONNECTION_PROFILE—SSL VPN user login group drop-down, a group alias within the connection profile. • CSCO_WEBVPN_MACRO1—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value1. Variable substitution via RADIUS is performed by VSA#223. • CSCO_WEBVPN_MACRO2—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value2. Variable substitution via RADIUS is performed by VSA#224. • CSCO_WEBVPN_MACROLIST1 and CSCO_WEBVPN_MACROLIST2—Statically configured bookmarks which can use arbitrarily-sized lists provided by LDAP attribute maps. These macros take the following three parameters: – Delimiter—Delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. – Index—Index is an administrator-supplied integer which specifies the number of the element in the list to select. The value can range between 1 and 128. – URL Encoding—URL Encoding is the choice to apply the LDAP string before it is substituted into the ASA device's request. You can select one of the following values: – None—No transformation occurs on the string value before sending to the backend server. – url-encode—Each parsed value is URL encoded, except for a list of reserved characters that make up the special characters in a URL. – url-encode-data—Each parsed value is transformed fully with URL encoding. – base64—Each parsed value is base 64 encoded. • CSCO_WEBVPN_PRIMARY_USERNAME—Primary user login ID when double authentication is enabled and login ID has primary login username. • CSCO_WEBVPN_SECONDARY_USERNAME—Secondary user login ID when double authentication is enabled.

Table 34-8 Add or Edit VDI Server Dialog Box (continued)

Element	Description
Password	Password for logging into the virtualization infrastructure server. This value can be a clientless macro. The macros available for password are: • CSCO_WEBVPN_PASSWORD—SSL VPN user login password. • CSCO_WEBVPN_INTERNAL_PASSWORD—SSL VPN user internal resource password. This is a cached credential, and not authenticated by a AAA server. If a user enters this value, it is used as the password for auto sign-on, instead of the password value. • CSCO_WEBVPN_MACRO1—Password for the MACRO1 username. • CSCO_WEBVPN_MACRO2—Password for the MACRO2 username. • CSCO_WEBVPN_MACROLIST1 and CSCO_WEBVPN_MACROLIST2—Statically configured bookmarks which can use arbitrarily-sized lists provided by LDAP attribute maps. These macros take the following three parameters: – Delimiter—Delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. – Index—Index is an administrator-supplied integer which specifies the number of the element in the list to select. The value can range between 1 and 128. – URL Encoding—URL Encoding is the choice to apply the LDAP string before it is substituted into the ASA device's request. You can select one of the following values: – None—No transformation occurs on the string value before sending to the backend server. – url-encode—Each parsed value is URL encoded, except for a list of reserved characters that make up the special characters in a URL. – url-encode-data—Each parsed value is transformed fully with URL encoding. – base64—Each parsed value is base 64 encoded. • CSCO_WEBVPN_PRIMARY_PASSWORD—Primary user login password for double authentication. • CSCO_WEBVPN_SECONDARY_PASSWORD—Secondaryuser login ID for double authentication.

ASA Group Policies SSL VPN Full Client Settings

Use the Full Client settings to configure the full client mode of access to the corporate network in a remote access SSL VPN for the ASA group policy object.

Full client mode enables access to the corporate network completely over an SSL VPN tunnel. In full client access mode, the tunnel connection is determined by the group policy configuration. The full client software, SSL VPN Client (SVC) or AnyConnect, is downloaded to the remote client, so that a tunnel connection is established when the remote user logs in to the SSL VPN gateway.

Tip To enable full client access, you must configure the **Remote Access VPN > SSL VPN > Other Settings** policy on the device to identify AnyConnect image packages to install on the device. The images must be on the device so that users can download them. For more information, see Understanding SSL VPN AnyConnect Client Settings, page 31-62 and Add and Edit File Object Dialog Boxes, page 34-37.

The following policies are supported for ASA 9.5(2) Remote Access VPN in Multi-context mode:

- Security Group Tag
- Periodic Certificate Verification
- Client Dead Peer Detection Timeout
- Gateway Dead Peer Detection Timeout
- Datalayer Transport layer Security Compression
- Keep AnyConnect Client on Client System
- Ignore Routing and Filter Rules
- AnyConnect Modules
- AnyConnect MTU
- AnyConnect Firewall-Client Public ACL
- AnyConnect Firewall-Client Private ACL
- Enable Datagram Transport Layer Security

Navigation Path

Select **SSL VPN > Full Client** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-9 ASA Group Policies SSL VPN Full Client Settings

Element	Description
Enable Full Client	Whether to enable full client mode.

Table 34-9 ASA Group Policies SSL VPN Full Client Settings (continued)

Element	Description
Mode	The mode in which to operate the SSL VPN: • **Use Other Access Modes if AnyConnect Client Download Fails**—If the full client fails to download to the remote user, allow the user to make clientless or thin client access to the VPN. • **Full Client Only**—Prohibit clientless or thin client access. The user must have the full client installed and functional to connect to the VPN.
Keep AnyConnect Client on Client System	Whether to leave the AnyConnect client installed on the client system after the client disconnects. If you do not leave the client installed, it must be download each time the user connects to the gateway.
Enable Keepalive Messages	Whether to exchange keepalive messages between peers to demonstrate that they are available to send and receive data in the tunnel. Keepalive messages transmit at set intervals, and any disruption in that interval results in the creation of a new tunnel using a backup device. If you select this option, enter the time interval (in seconds) that the remote client waits between sending IKE keepalive packets in the **Interval** field.
SSL Compression	Whether to enable data compression, and if so, the method of data compression to use: None, Deflate, or LZS. Data compression speeds up transmission rates for remote dial-in users connecting with modems. ⚠ **Caution** Data compression increases the memory requirement and CPU usage for each user session and consequently decreases the overall throughput of the security appliance. For this reason, it is recommended that you enable data compression only for remote users connecting with a modem. Design a group policy specific to modem users and enable compression only for them.
Client Dead Peer Detection Timeout (sec)	The time interval, in seconds, that the Dead Peer Detection (DPD) timer is reset each time a packet is received over the SSL VPN tunnel from the remote user. DPD is used to send keepalive messages between peer devices only when no incoming traffic is received and outbound traffic needs to be sent.
Gateway Dead Peer Detection Timeout (sec)	The time interval, in seconds, that the Dead Peer Detection (DPD) timer is reset each time a packet is received over the SSL VPN tunnel from the gateway.

Table 34-9 ASA Group Policies SSL VPN Full Client Settings (continued)

Element	Description
Key Renegotiation Method	The method by which the tunnel key is refreshed for the remote user group client: • **Disabled**—Disables the tunnel key refresh. • **Use Existing Tunnel**—Renegotiates the SSL tunnel connection. • **Create New Tunnel**—Initiates a new tunnel connection. Enter the time interval (in minutes) between the tunnel refresh cycles in the **Interval** field.
Enable Datagram Transport Layer Security	Whether to enable Datagram Transport Layer Security (DTLS) connections for the group. Enabling DTLS allows the AnyConnect client establishing an SSL VPN connection to use two simultaneous tunnels, an SSL tunnel and a DTLS tunnel. Using DTLS avoids latency and bandwidth problems associated with some SSL connections and improves the performance of real-time applications that are sensitive to packet delays.
Datagram Transport Layer Security Compression	Whether to compress Datagram Transport Layer Security (DTLS) connections for the group, and if so, the method of data compression to use: None, Default, or LZS.
Ignore Don't Fragment (DF) bit	Whether to ignore the DF bit in packets that need fragmentation. This feature allows the force fragmentation of packets that have the DF bit set, allowing them to pass through the tunnel. An example use case is for servers in your network that do not respond correctly to TCP MSS negotiations.

Table 34-9 ASA Group Policies SSL VPN Full Client Settings (continued)

Element	Description
AnyConnect Module	The modules that the AnyConnect client needs to enable optional features. Click **Select** to select the applicable modules from the Add AnyConnect Module dialog box. • **AnyConnect DART**—Select this module to enable the AnyConnect Diagnostics and Reporting Tool (DART), which bundles specified log files and diagnostic information that can be used for analyzing and debugging the client connection. • **AnyConnect Network Access Manager**—Select this module to enable the Network Access Manager, which enforces administratively defined end user and authentication policies and makes the pre-configured network profiles available to end users. • **AnyConnect SBL**—Select this module to enable Start Before Logon (SBL), which forces the user to connect to the enterprise infrastructure over a VPN connection before logging on to Windows by starting AnyConnect before the Windows login dialog box appears. After authenticating to the ASA, the Windows login dialog appears, and the user logs in as usual. SBL is only available for Windows and lets you control the use of login scripts, password caching, mapping network drives to local drives, and more. • **AnyConnect Web Security Module**—Select this module to enable the AnyConnect Web Security module, which is an endpoint component that routes HTTP traffic to a ScanSafe scanning proxy where the ScanSafe web scanning service evaluates it. • **AnyConnect Telemetry Module**—Select this module to enable the AnyConnect telemetry module for AnyConnect Secure Mobility Client, which sends information about the origin of malicious content to the web filtering infrastructure of the Cisco IronPort Web Security Appliance (WSA). The web filtering infrastructure uses this data to strengthen its web security scanning algorithms, improve the accuracy of the URL categories and web reputation database, and ultimately provide better URL filtering rules. • **AnyConnect ISE Network Setup Assistant**—Select this module to enable the AnyConnect ISE Network Setup Assistant module. • **AnyConnect ISE Posture**—Select this module to enable the AnyConnect ISE Posture module. • **AnyConnect Posture Module**—Select this module to enable the AnyConnect Posture Module, which provides the AnyConnect Secure Mobility Client the ability to identify the operating system, antivirus, antispyware, and firewall software installed on the host. The Host Scan application, which is among the components delivered by the posture module, is the application that gathers this information. **Note** If other options are listed, see the release notes for the Cisco AnyConnect VPN Client for an explanation of the feature.

Table 34-9 ASA Group Policies SSL VPN Full Client Settings (continued)

Element	Description
AnyConnect MTU	The maximum transmission unit (MTU) size for SSL VPN connections established by the Cisco AnyConnect VPN Client.
AnyConnect Always-On VPN	Always-On VPN enables AnyConnnect to automatically establish a VPN session after you log onto the system. Note that until you log off from the system, the VPN session will remain open. Select one of these options: • None—AnyConnect service profile remains unchanged. It inherits the value from the default group policy. • AnyConnect Profile Setting—Always-On VPN option configured in the AnyConnect VPN profile is used by the AnyConnect client. • Disable—disables the Always-On VPN option.
AnyConnect Profile Name	The name of the AnyConnect profile to use for the group. You can enter multiple profile names each separated by a comma. You must configure this name and relate it to a profile in the **Remote Access VPN > SSL VPN > Other Settings** policy. **Note** The AnyConnect Profile name is supported from Security Manager version 4.12 for ASA devices running version 9.6(2) in Multi-context mode. The supported CLIs are: – Webvpn—Anyconnect profiles
Prompt User to Choose Client Time User Has to Choose Default Location	Whether to ask the user to download the client. Enter the number of seconds the user has to make a selection in the **Time User Has to Choose** field. The default is 120 seconds. If you do not select this option, the user is immediately taken to the default location. The user is also taken to the default location after the time to choose expires. • **Web Portal**—The portal page is loaded in the web browser. • **AnyConnect Client**—The AnyConnect client is downloaded.
Security Group Tag	ASA Version 9.3(1)+ supports security group tagging of VPN sessions. A Security Group Tag (SGT) can be assigned to a VPN session using an external AAA server, or by configuration of the local user database. This tag can then be propagated through the Cisco TrustSec system over Layer 2 Ethernet. Security group tags are useful on group policies and for local users when the AAA server cannot provide an SGT. When the Default check box is selected, no Security Group Tag is assigned. To specify a Security Group Tag, clear the **Default** check box and then enter the numerical value of the SGT tag that will be assigned to VPN users connecting with this group policy in the Security Group Tag field. Valid values are from 2 to 65519.

Table 34-9 ASA Group Policies SSL VPN Full Client Settings (continued)

Element	Description
Periodic Certificate Verification	Whether to enable periodic validation and revocation checking of the client certificates in VPN sessions. If you select this option, enter the interval of time, in hours, between 1 to 168. This feature is supported only in devices running ASA software version 9.4(1) or higher. Periodic certificate verification is disabled by default.
AnyConnect Firewall-Client Public ACL	The name of the Extended or Unified access control list or policy object to use to restrict user access to the SSL VPN. Public rules are applied to all interfaces on the client. Enter the name of the object or click **Select** to select it from a list or to create a new object. Unified ACLs are supported from ASA version 9.0. The default is Extended. If the device version is higher than ASA 9.0, all the Anyconnect values are discovered as Unified ACL and deployed during deployment.
AnyConnect Firewall-Client Private ACL	The name of the Extended or Unified access control list policy object to use to restrict user access to the SSL VPN. Private rules are applied to the Virtual Adapter. Enter the name of the object or click **Select** to select it from a list or to create a new object. Unified ACLs are supported from ASA version 9.0. The default is Extended. If the device version is higher than ASA 9.0, all the Anyconnect values are discovered as Unified ACL and deployed during deployment.
AnyConnect Custom Attributes table	The AnyConnect Custom Attribute table lists the custom attributes, names, and the corresponding values that are assigned to this group policy. AnyConnect custom attributes that are defined on the AnyConnect Custom Attribute tab of the SSL VPN Other Settings page are listed here (see Configuring AnyConnect Custom Attributes (ASA), page 31-70). Beginning with version 4.7, Security Manager enables to add a Custom Attribute Data to an existing Custom Attribute Type. You can add or remove the custom attributes for a group policy, and configure values for each attribute. • To add a custom attribute and its value click the **Add Row** button beneath the table and fill in the Add AnyConnect Custom Attribute dialog box. • To edit a custom attribute and its value, select it, click the **Edit Row** button, and make your changes in the Edit AnyConnect Custom Attribute dialog box. • To delete a custom attribute, select it and click the **Delete Row** button. You are asked to confirm the deletion. For more details, see Add/Edit AnyConnect Custom Attribute Dialog Box, page 31-71.

ASA Group Policies SSL VPN Settings

Use the SSL VPN Settings to configure attributes that are required for clientless and port forwarding (thin client) access modes to work, including auto signon rules for user access to servers. Auto Signon configures the security appliance to automatically pass SSL VPN user login credentials (username and password) on to internal servers. You can configure multiple auto signon rules.

The Homepage URL policy is supported for the SSL tab in ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **SSL VPN > Settings** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-10 ASA Group Policies SSL VPN Settings

Element	Description
Home Page	The URL of the SSL VPN home page. The URL is free text. The page is displayed when users log into the VPN. If you do not enter a URL, no home page is displayed. Beginning with version 4.12, Security Manager supports IPv6 address in the Home Page URL for ASA devices running the software version 9.0 or later. The format for the Home Page URL for IPv6 address is: http://[IPv6 address]/appname.The Home page URL should be prefixed with http:// (or) https://
Authentication Failure Message	The message to deliver to a remote user who successfully logs into the VPN but has no VPN privileges, and so can do nothing. The default message is: "Login was successful, but because certain criteria have not been met or due to some specific group policy, you do not have permission to use any of the VPN features. Contact your IT administrator for more information."
Minimum Keepalive Object Size (kilobytes)	The minimum size (in kilobytes) of an IKE keepalive packet that can be stored in the cache on the security appliance.
Single Sign On Server	The name of the single sign on (SSO) server policy object that identifies the server to use for this group, if any. An SSO server allows users to enter their username and password once and be able to access other server in the network without logging into each of them. If configure an SSO server, also configure the auto signon rules table. Enter the name of the object or click **Select** to select it from a list or to create a new object. For more information, see Add or Edit Single Sign On Server Dialog Boxes, page 34-42.
Enable HTTP Compression	Whether to allow an HTTP compressed object to be cached on the security appliance.

Table 34-10 ASA Group Policies SSL VPN Settings (continued)

Element	Description
Auto Signon Rules table	If you configure a single sign on server, the auto signon rules table contains the rules that determine which internal servers are provided the user's credentials. Thus, you can provide single sign on for some servers in your network but not others. Each rule is an allow rule, and indicates the IP address, subnet, or Universal Resource Identifier (URI) that identifies the server, and the type of authentication that will be sent to the server when the user tries to access it (either basic HTML, NTLM, FTP, or all of these). The rules are processed in order, top to bottom, and the first match is applied. Therefore, be sure to order the rules correctly using the up and down arrow buttons. If the user accesses a server that is not identified in one of these rules, the user must log into the server to gain access. • To add a rule, click the **Add Row** button to open the Add or Edit Auto Signon Rules Dialog Box, page 34-27. • To edit a rule, select it and click the **Edit Row** button. • To delete a rule, select it and click the **Delete Row** button.
Portal Page Customization	The name of the SSL VPN customization policy object that defines the appearance of the portal web page. The portal page allows the remote user access to all the resources available on the SSL VPN network. If you do not specify an object, the default page appearance is used. Enter the name of the object or click **Select** to select it from a list or to create a new object. For more information, see Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78.
User Storage Location	The location where personalized user information is stored between clientless SSL VPN sessions. If you do not specify a location, information is not stored between sessions. Stored information is encrypted. Enter a file system designation in the following format: **protocol://username:password@host:port/path** Where **protocol** is the protocol of the server, **username** and **password** are a valid user account on the server, and **host** is the name of the server. Also indicate the **port** number (if you do not use the default for the protocol) and directory **path** of the location on the server to use. For example: **cifs://newuser:12345678@anyfiler02a/new_share**
Storage Key Confirm	The storage key used to protect data stored between sessions. Spaces are not supported.
Post Max Size	The maximum size allowed for a posted object. The range is 0 through 2147483647 (which is the default). Specify 0 to prevent posting.
Upload Max Size	The maximum size allowed for a uploaded object. The range is 0 through 2147483647 (which is the default). Specify 0 to prevent uploading.

Table 34-10 ASA Group Policies SSL VPN Settings (continued)

Element	Description
Download Max Size	The maximum size allowed for a downloaded object. The range is 0 through 2147483647 (which is the default). Specify 0 to prevent downloads.

Add or Edit Auto Signon Rules Dialog Box

Use the Add or Edit Auto Signon Rules dialog box to configure the Auto Signon rules that the security appliance uses to pass SSL VPN user login credentials on to an internal server.

Navigation Path

Open the ASA Group Policies SSL VPN Settings, page 34-25, then click **Create**, or select an item in the table and click **Edit**.

Field Reference

Table 34-11 Add or Edit Auto Signon Rules Dialog Box

Element	Description
Allow IP	Select this option to configure an IPv4 or IPv6 address or subnet for the rule. Any server within this subnet is supplied the specified login credentials. Beginning with version 4.12, Security Manager supports IPv6 addresses for devices running ASA 9.0 or later. • To enter the IP address of a single server, enter the full IP address and use 255.255.255.255 as the subnet mask. • To specify a subnet, enter the network address and subnet mask, for example, IP address 10.100.10.0 mask 255.255.255.0. If you want the appliance to send credentials to any internal server the user tries to access, create rules for all of your internal networks. You might be able to do this with a single rule.
Allow URI	Select this option to configure a Universal Resource Identifier (URI) for the rule. This identifies the internal server based on URI rather than IP address. For example, **https://*.example.com/*** creates a rule for all web pages on any server in the example.com domain. Use the asterisk as a wildcard to apply to zero or more characters.

Table 34-11 Add or Edit Auto Signon Rules Dialog Box (continued)

Element	Description
Login Credentials	Beginning with Security Manager version 4.7, you can select the login username and password from the available variables or macros. **Note** These macros are supported on devices running ASA software release version 8.2(1) and higher. The macros available for username are: • CSCO_WEBVPN_USERNAME—SSL VPN user login ID. • CSCO_WEBVPN_CONNECTION_PROFILE—SSL VPN user login group drop-down, a group alias within the connection profile. • CSCO_WEBVPN_MACRO1—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value1. Variable substitution via RADIUS is performed by VSA#223. • CSCO_WEBVPN_MACRO2—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value2. Variable substitution via RADIUS is performed by VSA#224. • CSCO_WEBVPN_MACROLIST1 and CSCO_WEBVPN_MACROLIST2—Statically configured bookmarks which can use arbitrarily-sized lists provided by LDAP attribute maps. These macros take the following three parameters: – Delimiter—Delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. – Index—Index is an administrator-supplied integer which specifies the number of the element in the list to select. The value can range between 1 and 128. – URL Encoding—URL Encoding is the choice to apply the LDAP string before it is substituted into the ASA device's request. You can select one of the following values: – None—No transformation occurs on the string value before sending to the backend server. – url-encode—Each parsed value is URL encoded, except for a list of reserved characters that make up the special characters in a URL. – url-encode-data—Each parsed value is transformed fully with URL encoding. – base64—Each parsed value is base 64 encoded. • CSCO_WEBVPN_PRIMARY_USERNAME—Primary user login ID when double authentication is enabled and login ID has primary login username. • CSCO_WEBVPN_SECONDARY_USERNAME—Secondary user login ID when double authentication is enabled. The macros available for password are: • CSCO_WEBVPN_PASSWORD—SSL VPN user login password.

Table 34-11 Add or Edit Auto Signon Rules Dialog Box (continued)

Element	Description
Authentication Type	The type of credentials that the security appliance will pass on to the servers covered by this rule: Basic HTML, NTLM (NT LAN Manager) authentication, FTP, or all of these methods. The default option is **All**. Use the default unless you want to limit logins to a certain type.

ASA Group Policies Browser Proxy Settings

Use the Browser Proxy settings to configure the attributes for the browser.

Browser Proxy is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Browser Proxy** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-12 ASA Group Policies Browser Proxy Settings

Element	Description
Proxy Server Policy	Select one of the following: • No proxy—If you select this option, proxy settings will not be used. • Do not modify client proxy—If you select this option, ASA will not modify the proxy setting on the endpoint device. • Use proxy—If you select this option, select one or more methods available in Select Proxy Method.
Select Proxy Method	Select one or more of the following: • Auto Detect—Select this option to enable the use of automatic proxy server detection in the client device's browser. • Use Proxy Server Setting Configured Below—Select this option and then specify the proxy server settings. • User Proxy Auto Configuration (PAC) configured below—Select this option to direct the browser to retrieve the HTTP proxy server setting from the proxy auto-configuration file URL.

Table 34-12 ASA Group Policies Browser Proxy Settings (continued)

Element	Description
Proxy Server Setting	Enter the following: • Server Address—Specify the IP address or name and the port of the browser server that is applied for the client device in the format 'ServerAddress:PortNumber'. To configure multiple proxy servers, separate the server addresses using a space. • Exception List—List the server names and IP addresses that you want to exclude from proxy server access. Enter the list of addresses that you do not want to have accessed through a proxy server. This list corresponds to the Exceptions list in the Proxy Settings dialog box in the browser. To configure multiple exception lists, separate the lists using a space, comma, or semicolon. • Bypass Server for Local Addresses—Configures the browser proxy local-bypass settings for a client device. Click **Yes** to enable local bypass or **No**to disable local bypass. Select **None** if you do not want to use this option. The default selected option is None.
Proxy Auto Configuration (PAC) URL	Specify the URL of the auto-configuration file. This file tells the browser where to look for proxy information.
Policy Lockdown	Select **Enable** to hide the Connections tab in the browser for the duration of an AnyConnect VPN session. Select **Disable** to leave the display of the Connections tab unchanged. Select **None** if you do not want to use this option. The default selected option is None.

ASA Group Policies DNS/WINS Settings

Use the DNS/WINS settings to define the DNS and WINS servers and the domain name that should be pushed to clients associated with the ASA group policy. These settings apply to Easy VPN and remote access IPSec and SSL VPN configurations.

DNS/WINS is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **DNS/WINS** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-13 ASA Group Policies DNS/WINS Settings

Element	Description
Primary IPv4 DNS Server	The IPv4 address of the primary DNS server for the group. Enter the IPv4 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object. Primary IPv4 DNS Server address is mandatory to be able to configure Secondary IPv4 DNS Server.
Secondary IPv4 DNS Server	The IPv4 address of the secondary DNS server for the group. Enter the IPv4 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.

Table 34-13 ASA Group Policies DNS/WINS Settings (continued)

Element	Description
Primary IPv6 DNS Server	The IPv6 address of the primary DNS server for the group. Enter the IPv6 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices 9.0 or later. Primary IPv6 DNS Server address is mandatory to be able to configure Secondary IPv6 DNS Server.
Secondary IPv6 DNS Server	The IPv6 address of the secondary DNS server for the group. Enter the IPv6 address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices 9.0 or later.
Primary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
DHCP Network Scope	The scope of the DHCP network for the group. Enter the IP network address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Default Domain	The default domain name for the group. The default, blank, is none.

ASA Group Policies Split Tunneling Settings

Use the Split Tunneling settings to configure a secure tunnel to the central site and simultaneous clear text tunnels to the Internet. These settings apply to Easy VPN and remote access IPSec and SSL VPN configurations.

Split tunneling lets a remote client conditionally direct packets over an IPsec or SSL VPN tunnel in encrypted form or to a network interface in clear text form. With split tunneling enabled, packets not bound for destinations on the other side of the tunnel do not have to be encrypted, sent across the tunnel, decrypted, and then routed to a final destination. The split tunneling policy is applied to specific networks.

Split Tunneling is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Tip For optimum security, we recommend that you not enable split tunneling.

Navigation Path

Select **Split Tunneling** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-14 ASA Group Policies Split Tunneling Settings

Element	Description
DNS Names	A list of domain names to be resolved through the split tunnel. All other names are resolved using the public DNS server. If you do not enter a list, the list is inherited from the default group policy. Separate multiple entries with spaces or commas. The entire string can be a maximum of 255 characters.
Send all DNS traffic through the tunnel	Whether the AnyConnect client should resolve all DNS addresses through the VPN tunnel (SSL or IPsec/IKEv2). If DNS resolution through the tunnel fails, the address remains unresolved and the AnyConnect client does not try to resolve the address through public DNS servers. If you do not select this option, the client sends DNS queries over the tunnel according to the split tunnel policy specified by the Tunnel Option setting.
Tunnel Option	The policy you want to enable for split tunneling: • Disabled—(Default) No traffic goes in the clear or to any other destination than the security appliance. Remote users reach networks through the corporate network and do not have access to local networks. • Tunnel Specified Traffic—Tunnel all traffic from or to the networks permitted in the network ACL. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. • Exclude Specified Traffic—Traffic goes in the clear from and to the networks permitted in the network ACL. This is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel. This option applies only to the Cisco VPN Client.
IPv6 Tunnel Option	Beginning with version 4.10, Security Manager provides support for IPv6 traffic for Split Tunneling from ASA version 9.0. The policy you want to enable for split tunneling: • Disabled—(Default) No traffic goes in the clear or to any other destination than the security appliance. Remote users reach networks through the corporate network and do not have access to local networks. • Tunnel Specified Traffic—Tunnel all traffic from or to the networks permitted in the network ACL. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. • Exclude Specified Traffic—Traffic goes in the clear from and to the networks permitted in the network ACL. This is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel. This option applies only to the Cisco VPN Client.

Table 34-14 ASA Group Policies Split Tunneling Settings (continued)

Element	Description
Networks	The name of a standard, extended, or unified access control list policy object that identifies the networks that require traffic to travel across the tunnel and those that do not require tunneling. Unified ACLs are supported from ASA version 9.0. How permit and deny are interpreted depends on your selection for **Tunnel Option**. Enter the name of the object, or click **Select** to select it from a list or to create a new object. If you do not specify an ACL, the network list is inherited from the default group policy.

ASA Group Policies Connection Settings

Use the Connection Settings to configure the connection characteristics for the ASA group policy, including access control and session timeouts. These settings are used for Easy VPN and remote access IPsec or SSL VPN sessions.

Connection Settings is supported for ASA 9.5(2) Remote Access VPN in Multi-context mode.

Navigation Path

Select **Connection Settings** from the table of contents in the ASA Group Policies Dialog Box, page 34-1.

Field Reference

Table 34-15 ASA Group Policies Connection Settings

Element	Description
Filter ACL	The name of the extended access control list (ACL) policy object to use for filtering traffic on the VPN connection. The ACL determines which traffic is permitted or denied. Enter the name of the object or click **Select** to select it from a list or to create a new object. Beginning with version 4.10 and ASA version 9.0, you can select from a list of Standard, Extended, or Unified ACL objects. This ACL does not apply to clientless SSL VPN connections.
Banner Text	The banner, or welcome text, to display on remote clients when they connect to the VPN. • Beginning with version 4.9, Security Manager supports up to 4000 characters for the banner text for ASA devices on version 9.5(1) or higher. • For ASA version lower than 9.5(1), Security Manager allows you to enter up to 500 characters for the banner text.
IPv4 Address Pools	Specifies the name of one or more IPv4 address pools to use for this group policy. Enter the names of the IPv4 address pool objects separated by a comma or click **Select** to select the objects from a list or to create a new objects.

Table 34-15 ASA Group Policies Connection Settings (continued)

Element	Description
IPv6 Address Pools	Specifies the name of one or more IPv6 address pools to use for this group policy. Enter the names of the IPv6 address pool objects separated by a comma or click **Select** to select the objects from a list or to create a new objects. Beginning with version 4.12, Security Manager supports IPv6 address pools for ASA devices 9.0 or later.
Access hours	The name of a time range policy object that specifies the times that users are allowed to access the VPN. If you do not specify a time range, users can access the VPN at all times. Specify a time range if you want to limit access to the network to certain hours, such as the typical work days and work hours for your organization. Enter the name of the object or click **Select** to select it from a list or to create a new object. For more information, see Configuring Time Range Objects, page 6-71.
Max Simultaneous Logins	The number of simultaneous logins a single user is allowed. Values are 0-2147483647. The default is 3. Specify 0 to disable logins and prevent user access.
Max Connection Time	The maximum amount of time a user is allowed to be connected to the VPN. Select one of the following: • Specified Connection time—Use the maximum time value that you enter. Values are 1-35791394 minutes. After the time is exceeded, the security appliance closes the connection. • Unlimited Connection time—The security appliance does not close connections based on connection time.
Idle Timeout	The amount of time a user is allowed to be connected to the VPN while the connection is idle, that is, there is no communication activity. Select one of the following: • Specified Timeout—Use the time out value you enter. Values are 1-4473924 minutes. When the idle time is exceeded, the security appliance closes the connection. The default is 30 minutes. • Unlimited Timeout—The security appliance does not close idle connections.
VLAN Mapping VLAN ID	The VLAN ID value can be between 1 and 4094 and must correspond to a VLAN interface on the ASA. The VLAN mapping feature on the ASA allows for traffic from VPN connections to be directed to a specified VLAN interface. Beginning with Cisco Security Manager version 4.10 and ASA version 9.5(1), you can assign IPv6 addresses to remote users. Beginning with Cisco Security Manager version 4.17, you can configure VLAN on ASA 9.9(2) or later multi-context devices.

Add or Edit Secure Desktop Configuration Dialog Box

Use the Add or Edit Cisco Secure Desktop Configuration dialog box to create, copy, and edit Cisco Secure Desktop Configuration objects for IOS routers. You can configure the settings required for Windows clients who are connecting from different location types, enable or restrict web browsing and file access for Windows CE clients, and configure the cache cleaner for Macintosh and Linux clients.

Cisco Secure Desktop (CSD) secures network endpoints by providing a reliable means of eliminating all traces of sensitive data by providing a single, secure location for session activity and removal on the client system.

This policy object uses the Secure Desktop Manager application to configure the settings. For an example of configuring settings, see *Cisco Secure Desktop on IOS Configuration Example Using SDM* at http://www.cisco.com/en/US/products/ps6496/products_configuration_example09186a008072aa7b.shtml. The first part of the configuration example explains setting up SDM, which you can ignore. Instead, look for the sections that describe setting up Windows locations midway through the example. The screen shots will help you identify when you are looking at CSD configuration.

Navigation Path

Select **Manage > Policy Objects**, then select **Cisco Secure Desktop (Router)** from the Object Type Selector. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- Creating Cisco Secure Desktop Configuration Objects, page 33-18
- Policy Object Manager, page 6-4

Field Reference

Table 34-16 Add or Edit Secure Desktop Configuration Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object (up to 1024 characters).

Table 34-16 Add or Edit Secure Desktop Configuration Dialog Box (continued)

Element	Description
Windows Location Settings	
Windows Locations	The names of the locations that you want to configure for Windows clients connecting from specific locations, such as Work, Home, or Insecure. When you create a location, an item for the location is added to the table of contents, where you can select the settings folders related to the location and configure its properties. The settings include a definition of how to determine if a client is connecting from that particular location. For each location you want to configure, enter its name in the **Location to Add** field and click **Add** to move it to the **Locations** list. You can reorder the locations using the Move Up/Move Down buttons. CSD checks locations in the order listed in this dialog box, and grants privileges to client PCs based on the first location definition they match. You can create a default location, such as Insecure, as the final location and configure the strictest security for it. For more information, see Creating Cisco Secure Desktop Configuration Objects, page 33-18.
Close all open browser windows after installation	Whether to close all the open browser windows after installing the Secure Desktop application.
VPN Feature Policy	Select the check boxes to enable these features if installation or location matching fails: • Web Browsing • File Access • Port Forwarding • Full Tunneling
Windows CE	
VPN Feature Policy	The Windows CE options enable you to configure a VPN feature policy to enable or restrict web browsing and remote server file access for remote clients running Microsoft Windows CE. You cannot configure locations for these clients.
Mac and Linux Cache Cleaner	
Launch Cleanup Upon Global Timeout	Whether to set a global timeout after which CSD launches the cache cleaner. Select a timeout (the default is 30 minutes), and select whether to allow the user to reset the timeout value.
Launch Cleanup Upon Exiting of Browser	Whether to start the cache cleaner when the user closes all web browser windows.
Enable Canceling of Cleaning	Whether to allow the remote user to cancel the cleaning of the cache.

Table 34-16 Add or Edit Secure Desktop Configuration Dialog Box (continued)

Element	Description
Secure Delete	The number of passes for CSD to perform a secure cleanup. The default is 1 pass. CSD encrypts and writes the cache to the remote client's disk. Upon termination of the Secure Desktop, CSD converts all bits occupied by the cache to all 0's, then to all 1's, and then to randomized 0's and 1's.
Enable Web Browsing if Mac or Linux Installation Fails	Whether to allow web browsing (but not other remote access features) if the cache cleaner installation fails.
VPN Feature Policy	Whether to allow web browsing, remote server file access, and port forwarding for Macintosh and Linux clients. Port forwarding permits the use of the Secure Desktop to connect a client application installed on the local PC to the TCP/IP port of a peer application on a remote server.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Add and Edit File Object Dialog Boxes

Use the Add and Edit File Object dialog boxes to create, copy, and edit file objects. File objects represent files that are used in device configurations, typically for remote access VPN policies and policy objects. Such files include Anyconnect client profile and image files, image (graphic) files, plug-in jar files, and Cisco Secure Desktop package files.

When you create a file object, Security Manager makes a copy of the file in its storage system. These files are backed up whenever you create a backup of the Security Manager database, and they are restored if you restore the database. When you deploy configurations that specify a file object, the associated file is download to the device in the appropriate directory.

After you create a file object, you typically should not edit it. If you need to replace the file, edit the file object to select the new file, or create a new file object. If the file is editable, you can edit the file object to identify the file's location in the file repository, and use the desired editor to open and edit the file outside of Security Manager. The file repository is the **CSCOpx\MDC\FileRepository** folder in the installation directory (typically, C:\Program Files). The files are organized in subfolders named for the file type.

For all file types except Image files, you can add a file from the Security Manager server or from the local Security Manager client by selecting the appropriate tab on the Choose a file dialog box. You cannot select files from a network server. You can control the ability to add files from the Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10.

Tip If you are copying a file to the Security Manager server so that it can be used in a file object, do not copy the file directly to the file repository.

When you delete a file object, the associated file is not deleted from the file repository.

Navigation Path

Select **Manage > Policy Objects**, then select **File Objects** from the Object Type Selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Configuring SSL VPN Browser Plug-ins (ASA), page 31-60
- Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9
- SSL VPN Customization Dialog Box—Informational Panel, page 34-57
- SSL VPN Customization Dialog Box—Title Panel, page 34-53

Field Reference

Table 34-17 Add and Edit File Object Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9. If you do not enter a name, the name of the file is used for the object name.
Description	An optional description of the object.
File Type	The type of file. If you create the object while configuring a policy, the correct file type is pre-selected. Options are: • Image—For graphic files. • Cisco Secure Desktop Package • Plug-In—For browser plug-in files. • AnyConnect Profile • AnyConnect Image • Hostscan Image

Table 34-17 Add and Edit File Object Dialog Boxes (continued)

Element	Description
File	The name and full path of the file. Click **Browse** to select the file. The following file types are managed using Image Manager. For more information, see Image Manager Supported Image Types, page 73-5. • Cisco Secure Desktop Package • Plug-In—For browser plug-in files. • AnyConnect Image • Hostscan Image For AnyConnect Profile and Image files, you can add a file from the Security Manager server. You cannot select files from a network server. **Tip** You can control the ability to add files from the Security Manager client from **Tools > Security Manager Administration > Customize Desktop**. For more information, see Customize Desktop Page, page 11-10. For file objects that you are editing, the path indicates the location in the Security Manager file repository.
File Name on Device	The file name you want to use when the file is downloaded to the device when you deploy policies. The default is to use the same file name as the original file. If the object was created by discovering policies from the device, this field uses the original name of the file as it existed on the device. This might not be the same name as it exists on the Security Manager server if the original name duplicated existing file names on the server.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

File Object — Choose a file Dialog Box

Use the File Object — Choose a file dialog box to select the file to use for the file object you are adding or editing. The available files are managed using Image Manager. For more information, see Image Manager Supported Image Types, page 73-5.

Navigation Path

Select **Manage > Policy Objects**, then select **File Objects** from the Object Type Selector. Add or Edit a file object and from the Add or Edit File Object dialog box, click **Browse** to open the File Object — Choose a file dialog box.

Related Topics

- Understanding and Managing SSL VPN Support Files, page 30-5
- Add and Edit File Object Dialog Boxes, page 34-37
- Configuring SSL VPN AnyConnect Client Settings (ASA), page 31-64
- Configuring SSL VPN Browser Plug-ins (ASA), page 31-60

- Configuring Cisco Secure Desktop Policies on ASA Devices, page 32-9
- SSL VPN Customization Dialog Box—Informational Panel, page 34-57
- SSL VPN Customization Dialog Box—Title Panel, page 34-53

Field Reference

Table 34-18 File Object — Choose a file Dialog Box

Element	Description
Image Repository	Lists the available files you can use for defining your file object. The available files are managed using Image Manager. For more information, see Image Manager Supported Image Types, page 73-5.
File selected	Shows the currently select file object.
Files of Type	Filters the list of files. Options are: **Note** You can only view all file objects or only objects filtered by the type of file object you are adding or editing. • Cisco Secure Desktop Package • Plug-In—For browser plug-in files. • AnyConnect Image • Hostscan Image • All

Add or Edit Port Forwarding List Dialog Boxes

Use the Port Forwarding List dialog box to create, copy and edit port forwarding list policy objects. You can create port forwarding list objects to use when you are configuring the thin client access mode for SSL VPN.

Port forwarding allows users to access applications (such as Telnet, e-mail, VNC, SSH, and Terminal services) inside the enterprise through an SSL VPN session. When port forwarding is enabled, the hosts file on the SSL VPN client is modified to map the application to the port number configured in the forwarding list. A port forwarding list object defines the mappings of port numbers on the remote client to the application's IP address and port behind the SSL VPN gateway.

Navigation Path

Select **Manage > Policy Objects**, then select **Port Forwarding List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- SSL VPN Access Modes, page 30-4
- ASA Group Policies SSL VPN Clientless Settings, page 34-12
- User Group Dialog Box—Thin Client Settings, page 34-84
- Create Group Policy Wizard—Clientless and Thin Client Access Modes Page, page 30-24
- Policy Object Manager, page 6-4

Field Reference

Table 34-19 Port Forwarding List Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Port Forwarding List table	The port forwarding entries that are defined in the object. The entries show the mapping of the local port to the remote server and port. • To add a mapping, click the **Add Row** button to open the Add or Edit A Port Forwarding Entry Dialog Box, page 34-41. • To edit a mapping, select it and click the **Edit Row** button. • To delete a mapping, select it and click the **Delete Row** button.
Include Port Forwarding Lists	The names of other port forwarding list objects to include in the object. Enter the name of the object or click **Select** to select it from a list or to create a new object. Separate multiple entries with commas. When you add other port forwarding lists, the entries from those lists are treated as if they were directly entered into this object, and the names of the included objects are not reflected in the device configuration commands during deployment.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit A Port Forwarding Entry Dialog Box

Use the Add or Edit A Port Forwarding Entry dialog boxes to create a new port forwarding list entry or edit an existing one.

Navigation Path

Go to the Add or Edit Port Forwarding List Dialog Boxes, page 34-40 and click the **Add Row** button or select an entry and click the **Edit Row** button beneath the Port Forwarding List table.

Field Reference

Table 34-20 Add or Edit A Port Forwarding Entry Dialog Box

Element	Description
Local TCP Port	The port number to which the local application is mapped (between 1 and 65535).
Remote Server IPv4/IPv6 Address Name	The IPv4 or IPv6 address or fully qualified domain name of the remote server. Select the type of entry and enter the IP address or name. Beginning with version 4.12, Security Manager supports IPv6 addresses for ASA devices running the software version 9.0 or later. For the IP address, you can enter the name of a network/host object that specifies the remote server's IP address, or click **Select** to select it from a list or to create a new object.
Remote TCP Port	The port number of the application for which port forwarding is configured (between 1 and 65535).
Description	A description of the port forwarding entry. This information is mandatory on Cisco IOS devices.

Add or Edit Single Sign On Server Dialog Boxes

Use the Add or Edit Single Sign On Server dialog box to create, copy, and edit single sign on (SSO) server objects for use with SSL VPNs (as configured in ASA group policy objects). For information on how to configure SSO servers in an ASA group policy, see ASA Group Policies SSL VPN Settings, page 34-25.

Single sign-on lets users access different secure services on different servers without entering a username and password more than once. In the authentication, the security appliance acts as a proxy for the SSL VPN user to the SSO server. You can configure this object to identify either a Computer Associates SiteMinder SSO server or a Security Assertion Markup Language (SAML) Browser Post Profile version 1.1 server.

The SSO mechanism starts as part of the AAA process or just after successful user authentication to an AAA server. The SSL VPN server running on the security appliance acts as a proxy for the user to the authenticating server. When a user logs in, the SSL VPN server sends an SSO authentication request, including username and password, to the authenticating server. If the server approves the authentication request, it returns an SSO authentication cookie to the SSL VPN server. The security appliance keeps this cookie on behalf of the user and uses it to authenticate the user to secure web sites within the domain protected by the SSO server.

If you want to configure SSO for an SSL VPN group, you must also configure a AAA server, such as a RADIUS or LDAP server.

Note The SAML Browser Artifact profile method of exchanging assertions is not supported.

Navigation Path

Select **Single Sign On Servers** in the Policy Object Manager, page 6-4. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

You can also create the object when configuring an ASA user group object for SSL VPN (see ASA Group Policies SSL VPN Settings, page 34-25).

Field Reference

Table 34-21 Add or Edit Single Sign-On Server Dialog Box

Element	Description
Name	The object name, which must be 4 to 31 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Authentication Type	The type of SSO server to use with clientless SSL VPN connections. The other attributes on the page change based on your selection. • **SiteMinder**—Computer Associates SiteMinder SSO server. • **SAML POST**—Security Assertion Markup Language (SAML) Browser Post Profile server.
URL (SiteMinder only.)	The URL of the SiteMinder SSO server to which the security appliance makes authentication requests. Select whether to use HTTP or HTTPS and enter the URL. **Tip** For HTTPS communication, make sure that the SSL encryption settings match on both the security appliance and the SiteMinder server. On the security appliance, you can verify this with the **ssl encryption** command.
Secret Key Confirm (SiteMinder only.)	The key used to encrypt authentication communications with the SiteMinder server, if any. The key can contain any alphanumeric characters. There is no minimum or maximum number of characters. Enter the same key in both fields. **Tip** If you enter a secret key, you must configure the same key in the SiteMinder server using the Cisco Java plug-in authentication scheme.
Assertion URL (SAML POST only.)	The URL for the SAML-type SSO assertion consumer service. Select whether to use HTTP or HTTPS and enter the URL, which must be fewer than 255 characters.
Assertion Issuer (SAML POST only.)	The name of the security device that is sending assertions to a SAML-type SSO server. This is usually the name of the security appliance, for example, asa.example.com. The name must be fewer than 65 characters.
Trustpoint (SAML POST only.)	The name of the PKI enrollment policy object that identifies the certificate authority (CA) server that acts as the trustpoint that contains the certificate to use to sign the SAML-type browser assertion. Enter the name or click **Select** to select it from a list or to create a new object.
Max Retries	The number of times the security appliance retries a failed SSO authentication attempt before the authentication times out. The range is 1 to 5 retries, and the default is 3 retries.
Request Timeout	The number of seconds before a failed SSO authentication attempt times out. The range is 1 to 30 seconds, and the default is 5 seconds.

Table 34-21 Add or Edit Single Sign-On Server Dialog Box (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Bookmarks Dialog Boxes

Use the Add or Edit Bookmarks dialog boxes to configure browser-based clientless SSL VPN bookmarks (URL lists) for an SSL VPN Bookmark object. From this dialog box, you can change the order of the bookmark entries within the table, create, copy, edit, and delete SSL VPN Bookmark objects.

An SSL VPN Bookmark object defines the URLs that are displayed on the portal page after a successful login.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Bookmarks** from the Object Type Selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84
- Localizing SSL VPN Web Pages for ASA Devices, page 31-80

Field Reference

Table 34-22 Add and Edit Bookmarks Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Bookmarks Heading (IOS) (IOS devices only)	The heading that is displayed above the URLs listed on the portal page of an SSL†VPN hosted on an IOS device.

Table 34-22 Add and Edit Bookmarks Dialog Boxes (continued)

Element	Description
Bookmarks	The list of bookmark entries for the object. • To change the order of an entry, select it and click the Move Up or Move Down arrow buttons. The order of entries in the table defines the order in which the bookmarks are presented to the user. • To add an entry, click the Add button and fill in the Add Bookmark Entry dialog box (see Add or Edit Bookmark Entry Dialog Boxes, page 34-45). • To edit an entry, select it and click the Edit button. • To delete an entry, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Bookmark Entry Dialog Boxes

Use the Add or Edit Bookmark Entry dialog boxes to create or edit a bookmark to be included in an SSL VPN Bookmark object.

You can use non-English, non-ASCII languages for the text to display for bookmarks if you are configuring the object for use on an ASA device. For more information about how you can configure the SSL VPN portal in local languages, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

Navigation Path

In the Policy Object Manager, from the Add or Edit Bookmarks Dialog Boxes, right-click inside the Bookmarks table, then select **Add Row** or right-click a row, then select **Edit Row**.

Related Topics

- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84

Field Reference

Table 34-23 Add or Edit Bookmark Entry Dialog Boxes

Element	Description
Bookmark Option	Select whether you want to define a new SSL VPN Bookmark entry or use the entries from an existing object: • **Enter Bookmark**—You want to define a bookmark entry. • **Include Existing Bookmarks**—You want to include bookmark entries defined in an existing SSL VPN Bookmark object. Enter the name of the object or click **Select** to select it from a list or to create a new object. • **Predefined Application Templates**—You want to use a predefined template that contains the pre-filled necessary values for certain well-defined applications.
Select Auto sign-on Application	If you selected Predefined Application Templates as the Bookmark Option, select the auto sign-on application whose template you want to use: • Citrix XenApp • Citrix XenDesktop • Domino Web Access • Microsoft Outlook Web Access 2010 • Microsoft Outlook Web Access 2013 (ASA 9.4(1)+ only) • Microsoft SharePoint 2007 • Microsoft SharePoint 2010 • Microsoft SharePoint 2013 (ASA 9.5(1)+ only) • Citrix StoreFront 2.1 (ASA 9.3(1)+ only) • Citrix StoreFront 2.5 (ASA 9.4(1)+ only) After selecting an auto sign-on application, the Advanced Form and URL Settings are populated based on the selected application.
Title	The text label that the user sees for the bookmark.
URL	The Universal Resource Locator address for the bookmark. Select the protocol for the bookmark and enter the rest of the URL in the edit box. **Tip** If you are creating bookmarks for use on an ASA device, and you are also configuring Kerberos Constrained Delegation on the device, you might need to add the service principle name (SPN) to the URL. For more information, see Configuring Kerberos Constrained Delegation (KCD) for SSL VPN (ASA), page 31-69.
Settings These settings are applicable only to SSL VPN portals hosted on ASA devices running software version 8.x or later. Do not configure these settings for SSL VPN Bookmark objects that you will use on other devices.	
Subtitle	An additional user-visible title that describes the bookmark entry.

Table 34-23 Add or Edit Bookmark Entry Dialog Boxes (continued)

Element	Description
Thumbnail	The File object that represents an icon you want to associate with the bookmark on the Portal. Enter the name of the File object or click **Select** to select it from a list or to create a new object.
Authentication Access	Whether to display the thumbnail only on the Portal page. If you deselect this option, the thumbnail is also displayed on the Logon page.
Enable Favorite URL Option	Whether to display the bookmark entry on the portal home page. Deselect the check box if you want the bookmark entry to appear on the application page only.
Advanced Form and URL Settings These settings are applicable only to SSL VPN portals hosted on ASA devices running software version 8.x or later. Do not configure these settings for SSL VPN Bookmark objects that you will use on other devices.	
URL Method	Select the required URL method from the list: • **Get**—Select this option if you want simple data retrieval. • **Post**—Select this option when processing the data might involve changes to it, for example, storing or updating data, ordering a product, or sending e-mail. If you select this option, you must configure the Post parameters in the Post Parameters table. • **Auto Sign-on Form**—Select this option if you want to use auto sign-on.
Enable Smart Tunnel Option (Get and Post URL Method only)	Whether to open the bookmark in a new window that uses the smart tunnel functionality to pass data to and from the security appliance.
Preload Page Options (Get and Post URL Method only)	Optionally, configure the following Preload options: **Preload URL**—The URL of a page to load before the bookmark link is loaded. **Wait Time**—The time to allow for loading of the page before you are forwarded to the actual POST URL.

Table 34-23 Add or Edit Bookmark Entry Dialog Boxes (continued)

Element	Description
Auto Sign-on (ASA 9.0.1+ only) (Auto Sign-on Form URL Method only)	When Auto Sign-on Form is selected as the URL Method, configure the following options: **Note** Wildcards can be used in the URLs you enter for the following fields. For example, you can enter http*://www.example.com/myurl*. **Login Page URL**—The URL of the login page for which to auto sign-on. **Landing Page URL**—The URL of the page that is loaded after a successful login. The ASA requires the Landing Page to be configured to detect a successful login to the application. **Pre-Login Page URL**—The URL of the page which is loaded before the login page. This page will require user interaction to proceed to the login screen. **Control ID**—The ID of the control/tag that will get a click event on the pre-login page URL to proceed to the login page.
Post Parameters	The list of the names and values of the Post parameters for the bookmark entry. • To add a parameter, click the Add button and fill in the Add Post Parameter dialog box (see Add and Edit Post Parameter Dialog Boxes, page 34-48). • To edit a parameter, select it and click the Edit button. • To delete a parameter, select it and click the Delete button.
Post Script	An optional field for entering JavaScript required by some applications. Some Web applications, such as Microsoft Outlook Web Access, may execute a JavaScript to change the request parameters before the log-on form is submitted.

Add and Edit Post Parameter Dialog Boxes

Use the Add and Edit Post Parameter dialog boxes to create a new Post parameter entry or edit an existing one in the table. For a detailed discussion of Post parameters, see Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84.

Navigation Path

In the Policy Object Manager, from the Add or Edit Bookmark Entry Dialog Boxes, right-click inside the Post Parameters table, then select **Add Row** or right-click a row, then select **Edit Row**.

Related Topics

- Configuring SSL VPN Bookmark Lists for ASA and IOS Devices, page 31-82
- Using the Post URL Method and Macro Substitutions in SSL VPN Bookmarks, page 31-84

Field Reference

Table 34-24 Add and Edit Post Parameter Dialog Boxes

Element	Description
Name	The name of the post parameter exactly as defined in the corresponding HTML form. For example, **param_name** in <input name="*param_name*" value="*param_value*">.

Table 34-24 Add and Edit Post Parameter Dialog Boxes (continued)

Element	Description
Value	The value of the post parameter exactly as defined in the corresponding HTML form. For example, **param_value** in <input name="*param_name*" value="*param_value*">. Select one of the following: • CSCO_WEBVPN_USERNAME—SSL VPN user login ID. • CSCO_WEBVPN_PASSWORD—SSL VPN user login password. • CSCO_WEBVPN_INTERNAL_PASSWORD—SSL VPN user internal resource password. This is a cached credential, and not authenticated by a AAA server. If a user enters this value, it is used as the password for auto sign-on, instead of the password value. • CSCO_WEBVPN_CONNECTION_PROFILE—SSL VPN user login group drop-down, a group alias within the connection profile. • CSCO_WEBVPN_DYNAMIC_URL1—A single bookmark that can generate multiple bookmark links on the user's portal. This macro takes *delimiter* as an option, where delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. • CSCO_WEBVPN_DYNAMIC_URL2—A single bookmark that can generate multiple bookmark links on the user's portal. This macro takes *delimiter* as an option, where delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. • CSCO_WEBVPN_MACRO1—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value1. Variable substitution via RADIUS is performed by VSA#223. • CSCO_WEBVPN_MACRO2—Set via the RADIUS/LDAP vendor-specific attribute. If you are mapping this from LDAP via an ldap-attribute-map, the Cisco attribute that uses this variable is WEBVPN-Macro-Substitution-Value2. Variable substitution via RADIUS is performed by VSA#224. • CSCO_WEBVPN_MACROLIST1 and CSCO_WEBVPN_MACROLIST2—Statically configured bookmarks which can use arbitrarily-sized lists provided by LDAP attribute maps. These macros take the following three parameters: – Delimiter—Delimiter is an administrator-supplied string which includes the characters used to separate the LDAP-mapped string into a list of values, using one delimiter per use of the macro. – Index—Index is an administrator-supplied integer which specifies the number of the element in the list to select. The value can range between 1 and 128. – URL Encoding—URL Encoding is the choice to apply the LDAP string before it is substituted into the ASA device's request. You can select one of the following values: – None—No transformation occurs on the string value before

Add and Edit SSL VPN Customization Dialog Boxes

Use the Add and Edit SSL VPN Customization dialog boxes to create, copy, and edit SSL VPN Customization objects. An SSL VPN Customization policy object describes how to customize web pages for a browser-based clientless SSL VPN hosted on an ASA 8.x device. For more information, see Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78.

You can use non-English, non-ASCII languages for the text to display on these pages. For more information about how you can configure the SSL VPN portal in local languages, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Customization** from the Object Type Selector. Right-click inside the work area, then select **New Object** or right-click a row, then select **Edit Object**.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78
- Localizing SSL VPN Web Pages for ASA Devices, page 31-80
- Creating Your Own SSL VPN Logon Page for ASA Devices, page 31-82

Field Reference

Table 34-25 Add and Edit SSL VPN Customization Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Settings Pane

The body of the dialog box is a pane with a table of contents on the left and settings related to the item selected in the table of contents on the right. Before configuring settings, click the Preview button to see the default settings to help you determine what, if anything, you want to change.

The top folders in the table of contents represent the SSL VPN web pages that you can customize, and are explained next.

Table 34-25 Add and Edit SSL VPN Customization Dialog Boxes (continued)

Element	Description
Logon Page	The Logon web page is the one users see first when connecting to the SSL VPN portal. It is used for logging into the VPN. Select the following items in the Logon Page folder in the table of contents to view and change the settings: • **Logon Page**—The Browser Window Title field defines the title of the web page, which is displayed in the browser's title bar. • **Title Panel**—The title displayed in the web page itself. For more information about the settings, see SSL VPN Customization Dialog Box—Title Panel, page 34-53. • **Language**—The languages you will support for the Logon, Portal, and Logout pages. For more information about the settings, see SSL VPN Customization Dialog Box—Language, page 34-54. • **Logon Form**—The labels and colors used in the form that accepts user logon information. For more information about the settings, see SSL VPN Customization Dialog Box—Logon Form, page 34-56. • **Informational Panel**—An extra informational panel for conveying information to users. For more information about the settings, see SSL VPN Customization Dialog Box—Informational Panel, page 34-57. • **Copyright Panel**—The copyright information on the logon page. For more information about the settings, see SSL VPN Customization Dialog Box—Copyright Panel, page 34-58. • **Full Customization**—If you do not want to use the security appliance's built-in logon page, even customized, you can instead enable full customization and supply your own web page. For more information about creating a custom Logon page and the settings, see Creating Your Own SSL VPN Logon Page for ASA Devices, page 31-82 and SSL VPN Customization Dialog Box—Full Customization, page 34-59.

Table 34-25 Add and Edit SSL VPN Customization Dialog Boxes (continued)

Element	Description
Portal Page	The Portal web page is the one users see after logging into the SSL VPN; it is the home page. Select the following items in the Portal Page folder in the table of contents to view and change the settings: • **Portal Page**—The **Browser Window Title** field defines the title of the web page, which is displayed in the browser's title bar. • **Title Panel**—The title displayed in the web page itself. For more information about the settings, see SSL VPN Customization Dialog Box—Title Panel, page 34-53. • **Toolbar**—The toolbar displayed above the main part of the Portal page. For more information about the settings, see SSL VPN Customization Dialog Box—Toolbar, page 34-59. • **Applications**—The application buttons that will appear on the page. For more information about the settings, see SSL VPN Customization Dialog Box—Applications, page 34-60. • **Custom Panes**—The layout of the main part of the Portal page. The default is a single column with no internal panes. For more information about the settings, see SSL VPN Customization Dialog Box—Custom Panes, page 34-61. • **Home Page**—How and whether to display URL lists on the home page. For more information about the settings, see SSL VPN Customization Dialog Box—Home Page, page 34-62.
Logout Page	The Logout web page is the one users see after logging out of the SSL VPN. For more information about the settings, see SSL VPN Customization Dialog Box—Logout Page, page 34-63.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

SSL VPN Customization Dialog Box—Title Panel

Use the Title Panel page of the SSL VPN Customization dialog box to determine whether the Logon page or Portal page will have a title displayed in the web page itself. If you enable the title panel, you can specify the title, font, font size and weight, styles, and colors used. You can also select a File object that identifies a logo graphic.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Title Panel** in the table of contents to configure the title of the Logon page, or **Portal Page > Title Panel** to configure the title of the Portal page.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78
- Localizing SSL VPN Web Pages for ASA Devices, page 31-80

Field Reference

Table 34-26 SSL VPN Customization Dialog Box—Title Panel

Element	Description
Display Title Panel	Whether to display a title panel within the web page. The default is to not display a title. If you select this option, you can configure the title using the other fields on this page.
Gradient	Whether to have the background color change in a gradual progression.
Title Text	The text to display in the title panel.
Font Weight Font Size Font Color	The characteristics of the font used for the title text. You can select a weight, font size, and color. Click **Select** to choose a font color.
Background Color	The color of the background of the title panel. Click **Select** to choose a color.
Style (CSS)	Cascading Style Sheet (CSS) parameters that define the style characteristics of the title panel. You can include a maximum of 256 characters. **Tip** For more information about CSS, visit the World Wide Web Consortium (W3C) website at www.w3.org.
Logo Image	The File policy object that identifies the logo image you want to include in the title panel, if any. Enter the name of the File object or click **Select** to select it from a list or to create a new object. **Tip** The image file can be a GIF, JPG, or PNG file, and it can be up to 100 kilobytes in size. For more information about File objects, see Add and Edit File Object Dialog Boxes, page 34-37.

SSL VPN Customization Dialog Box—Language

Use the Language page of the SSL VPN Customization dialog box identify the languages you will support on the browser-based clientless SSL VPN portal. If you want to configure translation tables for other languages on the ASA device and use them, you can configure the supported languages and allow users to choose their language. Before you configure these settings, read Localizing SSL VPN Web Pages for ASA Devices, page 31-80.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Language** in the table of contents.

Related Topics

- Localizing SSL VPN Web Pages for ASA Devices, page 31-80
- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-27 SSL VPN Customization Dialog Box—Language

Element	Description
Automatic Browser Language Selection	This table lists the languages you will support on the web pages for automatic browser language selection. Automatic browser language select allows the ASA device to negotiate with the user's web browser to determine the language in which to present the web pages. You must configure a translation table on the ASA device for any language you list here. For more detailed information about automatic browser language selection, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80. Languages are listed by their abbreviation in the table. The languages are evaluated top to bottom until a match is found. The language that is indicated as the default language (indicated as True in the table) is used if the device is unable to negotiate a different language with the browser. If you do not specify a default, English is the default. • To add a language, click the Add Row button below the table. • To edit a language, select it and click the Edit Row button. • To delete a language, select it and click the Delete Row button.
Enable Language Selector	Whether to display the Language Selector on the Logon page. The Language Selector allows users to select their preferred language. The Language Selector is complementary to the automatic browser language selection capability.
Language Selector Prompt	The text label for the Language Selector prompt.
Language Table	The list of languages included in the Language Selector drop-down list. You must configure a translation table on the ASA device for any language you list here. For more detailed information, see Localizing SSL VPN Web Pages for ASA Devices, page 31-80. The table lists the languages by abbreviation and title, or the common name of the language. The title is the text displayed in the drop-down list. You can change the language title but not the abbreviation. • To add a language, click the Add Row button below the table. • To edit a language, select it and click the Edit Row button. • To delete a language, select it and click the Delete Row button.

Add and Edit Language Dialog Boxes

Use the Add and Edit Language dialog boxes to add or edit an entry for a language you will support for automatic browser language selection or in the Language Selector drop-down list.

Navigation Path

From the SSL VPN Customization Dialog Box—Language page, click the **Add Row** button for either the Automatic Browser Language Selection table or the Language Selector table, or select a row and click the **Edit Row** button.

Related Topics

- Localizing SSL VPN Web Pages for ASA Devices, page 31-80
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-28 Add and Edit Language Dialog Boxes

Element	Description
Language	The list of languages that you can support on the browser-based clientless SSL VPN web pages, listed by their abbreviation.
Default (Automatic Browser Language Selection only)	Whether the language should be defined as the default language for the portal. The default language is used if the ASA device cannot negotiate a language with the client's browser.
Title (Language Selector only)	The name of the language that should appear in the Language Selector on the Logon page.

SSL VPN Customization Dialog Box—Logon Form

Use the Logon Form settings of the SSL VPN Customization dialog box to customize the title of the login box, login prompts of the SSL VPN page (including username, password, and group prompts), login buttons, and style elements of the login box that appears to browser-based clientless SSL VPN users when they initially connect to the security appliance.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Logon Form** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-29 SSL VPN Customization Dialog Box—Logon Page

Element	Description
Title	The text displayed as the title of the login box.

Table 34-29 SSL VPN Customization Dialog Box—Logon Page (continued)

Element	Description
Message	The message that appears in the login box above the username and password fields. You can enter a maximum of 256 characters.
Username Prompt	The text of the prompt for the username entry field.
Password Prompt	The text of the prompt for the password entry field.
Secondary Username Prompt Secondary Password Prompt	The prompts for a second username and password if you require two login credentials. You can enable secondary authentication only if the Connection Profile policy is configured to require it. The secondary username and password prompt are displayed only if you configure them. If you leave the username prompt blank, the primary username is used and the secondary password must be associated with the primary username.
Internal Password Prompt	The text of the prompt for the internal password entry field.
Show Internal Password First	Whether the prompt for the internal password should be placed above the password prompt. The internal password is required when using a clientless SSL VPN to access an internal protected website.
Group Selector Prompt	The text of the prompt for the Group Selector drop-down list.
Button Text	The name of the button the user clicks to log onto the SSL VPN.
Border Color	The color of the border of the login box. Click **Select** to choose a color.
Title Font Color	The color of the font for the login box title. Click **Select** to choose a color.
Title Background Color	The background color for the Title area of the login box. Click **Select** to choose a color.
Font Color	The color of the font of the login form. Click **Select** to choose a color.
Background Color	The background color for the login form. Click **Select** to choose a color.

SSL VPN Customization Dialog Box—Informational Panel

Use the Informational Panel page of the SSL VPN Customization dialog box to customize the appearance of the Informational panel in the Logon page. The Informational panel is an area where you can provide extra information to the user, and is optional.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Informational Panel** in the table of contents.

Related Topics

- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-30 SSL VPN Customization Dialog Box—Informational Panel

Element	Description
Display Informational Panel	Whether to display the Informational panel. The default is to not display the panel. If you select this option, you can configure the panel using the other fields on this page.
Panel Position	The location of the Informational panel, either to the left of the Logon box or to the right of it.
Text	The text that appears in the Informational panel. You can enter a maximum of 256 characters.
Logo Image	The File policy object that identifies the logo image you want to include in the Informational panel, if any. Enter the name of the File object or click **Select** to select it from a list or to create a new object. **Tip** The image file can be a GIF, JPG, or PNG file, and it can be up to 100 kilobytes in size. For more information about File objects, see Add and Edit File Object Dialog Boxes, page 34-37.
Image Position	The position of the logo image in the panel, either above the text or below it.

SSL VPN Customization Dialog Box—Copyright Panel

Use the Copyright Panel page of the SSL VPN Customization dialog box to customize the appearance of the Copyright panel in the Logon page. The Copyright panel provides your copyright information, appears at the bottom of the page, and is optional.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Copyright Panel** in the table of contents.

Related Topics

- Add and Edit SSL VPN Customization Dialog Boxes, page 34-51
- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-31 SSL VPN Customization Dialog Box—Copyright Panel

Element	Description
Display Copyright Panel	Whether to display the Copyright panel. The default is to not display the panel. If you select this option, you can configure the panel using the other fields on this page.
Text	The text that appears in the copyright panel. You can enter a maximum of 256 characters.

SSL VPN Customization Dialog Box—Full Customization

Use the Full Customization page of the SSL VPN Customization dialog box to identify your own custom Logon page. The custom page replaces the Logon page settings available on the dialog box. For information on creating a custom Logon page, see Creating Your Own SSL VPN Logon Page for ASA Devices, page 31-82.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logon Page > Full Customization** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-32 SSL VPN Customization Dialog Box—Full Customization

Element	Description
Enable Full Customization	Whether you want to use your own custom Logon page. If you enable full customization, all of the other Logon page configuration settings are ignored.
Custom Page	The custom Logon page. You must copy the file to the Security Manager server before specifying it here. Click **Browse** to select the file. For information on selecting files, see Selecting or Specifying a File or Directory in Security Manager, page 1-50.

SSL VPN Customization Dialog Box—Toolbar

Use the Toolbar page of the SSL VPN Customization dialog box to customize the appearance of the toolbar in the Portal page. The toolbar appears above the main body of the Portal page and includes a field to allow users to enter URLs to browse. The toolbar is optional.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Portal Page > Toolbar** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-33 SSL VPN Customization Dialog Box—Toolbar

Element	Description
Display Toolbar	Whether to display the toolbar. The default is to not display the toolbar. If you select this option, you can configure the toolbar using the other fields on this page.

Table 34-33 SSL VPN Customization Dialog Box—Toolbar (continued)

Element	Description
Prompt Box Title	The text of the prompt for the field where users select the protocol of the target web page and enter the URL.
Browse Button Text	The name of the button the user clicks to go to the target URL.
Logout Prompt	The text of the prompt for logging out of the SSL VPN.
User Prompt (only for ASA 9.7.1+)	The text of the prompt for a user currently logging into the remote access VPN.

SSL VPN Customization Dialog Box—Applications

Use the Applications page of the SSL VPN Customization dialog box to customize the application links that appear in the Portal page. This page lists all the application links that you can display in the navigational panel on the left side of the SSL VPN portal page.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Portal Page > Applications** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-34 SSL VPN Customization Dialog Box—Applications

Element	Description
No. Move Up and Move Down buttons (below the table)	The sequential number of the application in the table. To change the order of an application, select it and click the Move Up or Move down buttons to the desired position. The applications appear on the Portal page in the order represented here.
Application	The graphic associated with an application.
Title	The name of the application. Standard applications include Home, Web Applications, Browse Networks, Application Access, and AnyConnect Client. Also listed are the browser plug-ins that you create when you configure the SSL VPN global settings are also available for selection from this page. Double-click a title to make it editable so that you can change the name.
Enable	Whether the application is included on the Portal page.
Show Navigation Panel	Whether to display the navigation panel in the portal page. If you deselect this option, the list of applications does not appear on the portal.

SSL VPN Customization Dialog Box—Custom Panes

Use the Custom Panes page of the SSL VPN Customization dialog box to customize the appearance of the main body of the Portal page. By creating custom panes and specifying a column layout, you can create a grid of information that can help you present portal information effectively to your end users.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Portal Page > Custom Panes** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-35 SSL VPN Customization Dialog Box—Custom Panes

Element	Description
Columns table	The list of columns that the main body of the Portal page should be divided into. You define the column based on a percentage of the width of the page. The percentages should add up to 100. If they do not add up to 100, the device will adjust the column widths. Create the columns as you want them to appear, left to right, on the Portal page. • To add a column, click the Add Row button below the table. • To edit a column, select it and click the Edit Row button. • To delete a column, select it and click the Delete Row button.
Custom Panes table	The custom panes that should appear in the main body of the Portal page. The table shows whether a pane is enabled to appear, the type of pane, its characteristics, and the column and row in which it will appear on the page. The panes can display plain text or include a URL for HTML, image, or RSS links. For more detailed information about the settings, see Add or Edit Custom Pane Dialog Boxes, page 34-62. • To add a custom pane, click the Add Row button below the table. • To edit a custom pane, select it and click the Edit Row button. • To delete a custom pane, select it and click the Delete Row button.

Add and Edit Column Dialog Boxes

Use the Add or Edit Column dialog box to create or edit columns in the main body of the Portal page for browser-based clientless SSL VPNs. Enter the desired width of the column as a percentage of the total area in the Percentage field.

Navigation Path

From the SSL VPN Customization Dialog Box—Custom Panes page, click the **Add Row** button in the Column table, or select a column and click the **Edit Row** button.

Add or Edit Custom Pane Dialog Boxes

Use the Add or Edit Custom Pane dialog box to create or edit a pane to display in the main body or the Portal page of a browser-based clientless SSL VPN.

Navigation Path

From the SSL VPN Customization Dialog Box—Custom Panes page, click the **Add Row** button in the Custom Pane table, or select a pane and click the **Edit Row** button.

Field Reference

Table 34-36 Add and Edit Custom Pane Dialog Boxes

Element	Description
Enable	Whether to display the custom pane on the Portal page.
Type	The type of content to show in the pane, one of: • Text—Plain text. You can include HTML mark up. • HTML—HTML content provided by a URL. • Image—An Image provided by a URL. • RSS—An RSS feed provided by a URL.
Show Title Title	Whether to display a title in the pane. If you select this option, enter the title in the Title field.
Show Border	Whether to display a border around the pane.
Column Row	The column and row numbers in which the pane should appear. Select or enter the number for each to specify the desired grid location.
Height	The height of the pane in pixels.
URL (HTML, Image, and RSS content only.)	The URL that hosts the content you want to display in the pane.
Text (Text content only.)	The text you want to display in the pane. You can include HTML markup in the text.

SSL VPN Customization Dialog Box—Home Page

Use the Home Page page in the SSL VPN Customization dialog box to customize the appearance of the URL and file lists on the Portal page and the content of the main body of the Portal page. URL lists are considered to be default elements on the portal home page unless they are explicitly disabled.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Portal Page > Home Page** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-37 SSL VPN Customization Dialog Box—Home Page

Element	Description
Enable Custom Intranet Web Page	Whether to display a custom Intranet web page, which also enables URL bookmarks to be displayed on the Portal page. If you select this option, you can configure the panel using the other fields on this page.
URL List Mode	How you want to display URL lists on the home page. If you display URL lists, they are displayed in the column cells that are not occupied by custom panes (as configured on Portal Page > Custom Panes). The options are: • **Group By Application**—Bookmarks are grouped by application type. For example, Web Bookmarks, File Bookmarks. • **No Group**—URL lists are shown as separate panes. • **Do Not Display**—URL lists are not shown.
Custom Intranet Web Page URL	The URL of the custom web page that you want to be loaded as the home page. This page is displayed in the main body of the Portal page. If you specify a custom page, the settings on the Custom Panes page are ignored, and bookmark lists appear on the application pages that are accessed through the navigation panel on the left of the Portal page.

SSL VPN Customization Dialog Box—Logout Page

Use the Logout Page page of the SSL VPN Customization dialog box to customize the appearance of the Logout page for browser-based clientless SSL VPNs. The Logout page appears after the user logs out of the VPN.

Navigation Path

From the Add and Edit SSL VPN Customization Dialog Boxes, select **Logout Page** in the table of contents.

Related Topics

- Configuring ASA Portal Appearance Using SSL VPN Customization Objects, page 31-78

Field Reference

Table 34-38 SSL VPN Customization Dialog Box—Logout Page

Element	Description
Title	The text to display in the title panel.
Text	The message to display on the Logout page. Click Preview to see the default logout message. You can enter a maximum of 256 characters.
Show Login Button Login Button Text	Whether to display the Login button on the page. Displaying the button makes it easier for the user to log back into the portal. If you enable the button, you can specify the name of the button in the Login Button Text field.

Table 34-38 SSL VPN Customization Dialog Box—Logout Page (continued)

Element	Description
Border Color	The color of the border around the logout box. Click **Select** to choose a color.
Title Font Color Title Background Color	The color of the font and background for the title area of the page. Click **Select** to choose a color.
Font Color Background Color	The font and background color of the message that appears in the logout box. Click **Select** to choose a color.

Add or Edit SSL VPN Gateway Dialog Box

Use the Add or Edit SSL VPN Gateway dialog box to create, copy and edit SSL VPN gateway objects. You use these objects when you are configuring an SSL VPN connection on an IOS device. For more information, see SSL VPN Configuration Wizard—Gateway and Context Page (IOS), page 30-33.

An SSL VPN gateway acts as a proxy for connections to protected resources that are accessed through an SSL-encrypted connection between the gateway and a web-enabled browser on a remote device. You can configure only one gateway per SSL VPN.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Gateway** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- SSL VPN Configuration Wizard—Gateway and Context Page (IOS), page 30-33
- General Tab, page 33-16
- Policy Object Manager, page 6-4

Field Reference

Table 34-39 Add and Edit SSL VPN Gateway Dialog Boxes

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object (up to 1024 characters).

Table 34-39 Add and Edit SSL VPN Gateway Dialog Boxes (continued)

Element	Description
IP Address	The IP address for the gateway, which is the address to which remote users connect: • **Use Static IP Address**—Specify the address that you want to use. You must also configure this address on an interface on the router. • **Obtained from Interface**—Specify the interface role that resolves to a single interface on the device. The IP address configured for the interface is used. This option allows you to identify the external interface you want to use for connections without having to explicitly enter the IP address. If you have to change the address on the interface, you do not have to also reconfigure this object.
Port	The number of the port that will carry the HTTPS traffic. You can also enter the name of a port list object that specifies the single port number, or click **Select** to select the object from a list. The default is the HTTPS object, which specifies port 443. If you do not use port 443, you can enter another port number between 1025 and 65535.
Trustpoint	The digital certificate required to establish the secure connection. A self-signed certificate is generated when an SSL VPN gateway is activated.
Enable Gateway	Whether to activate the SSL VPN gateway.
Specify SSL Encryption Algorithms	Whether to restrict the encryption algorithms used for the connection, or to specify a different order of use. The default is to make all algorithms available in this order of preference: 3DES and SHA1, AES and SHA1, RC4 and MD5. Select the priority order for the algorithms. Select None to eliminate one or two algorithms.
Redirect HTTP Traffic HTTP Port	Whether to have the gateway redirect HTTP traffic over secure HTTP (HTTPS). Traffic that comes to this port is redirected to the port you specify in the Port field. Enter the port number for HTTP traffic in the **HTTP Port** field. You can enter a number or the name of a port list object, or click **Select** to select an object from a list or to create a new object. The HTTP port is normally 80. However, you can enter any other number that is used in your network between 1025-65535.
Hostname	The hostname for the gateway. • Do Not Specify—No hostname is assigned; the IP address to the gateway is used. • Use the host and domain names of the device—These are defined in the **Platform > Device Admin > Hostname** policy. • Use the Object—The hostname is the value defined in a text policy object. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 34-39 Add and Edit SSL VPN Gateway Dialog Boxes (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit Smart Tunnel List Dialog Boxes

Use the Add and Edit Smart Tunnel Lists dialog boxes to create, copy, and edit SSL VPN smart tunnel objects.

An SSL VPN smart tunnel list object lists the applications that are eligible for smart tunnel access to a private site. You can configure the clientless settings of an ASA group policy with a smart tunnel list to allow users to access the specified applications through the SSL VPN portal. For an explanation of the types of applications that support smart tunnel access, see Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85.

You can include other SSL VPN smart tunnel list objects in an object. Thus, you can create a smaller set of objects that identify your basic list of applications, then create other objects that create the required combination of applications. For example, you might want all three of your ASA group policies to allow smart tunnel access to applications A and B, but the remaining applications are unique for each group. By creating a single object that specifies A and B, you can include that object in each of the SSL VPN smart tunnel list objects for the group policies, and these objects need only specify their unique applications in the applications table.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Smart Tunnel Lists** from the Object Type selector. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- ASA Group Policies SSL VPN Clientless Settings, page 34-12
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Policy Object Manager, page 6-4

Field Reference

Table 34-40 Add and Edit Smart Tunnel Lists Dialog Boxes

Element	Description
Name	The object name, which can be up to 64 characters. Spaces are not allowed. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 34-40 Add and Edit Smart Tunnel Lists Dialog Boxes (continued)

Element	Description
Smart Tunnel Entries table	The applications to which users will be allowed smart tunnel access through the SSL VPN, including the name of the application and its location on client workstations. • To add an application, click the **Add Row** button to open the Add and Edit A Smart Tunnel Entry Dialog Boxes, page 34-67. • To edit an application, select it and click the **Edit Row** button. • To delete an application, select it and click the **Delete Row** button.
Include Smart Tunnel Lists	The other SSL VPN smart tunnel list objects that you want to include in this object, if any. Enter the names of the objects or click **Select** to select them from a list or to create new objects. Separate multiple entries with commas.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit A Smart Tunnel Entry Dialog Boxes

Use the Add and Edit A Smart Tunnel Entry dialog boxes to create a new smart tunnel entry or edit an existing entry in the table in the SSL VPN Smart Tunnel Lists dialog box.

Navigation Path

From Add and Edit Smart Tunnel List Dialog Boxes, page 34-66, click the **Add Row** button beneath the Smart Tunnel Entries table, or select an entry and click the **Edit Row** button.

Related Topics

- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Policy Object Manager, page 6-4

Field Reference

Table 34-41 Add and Edit Smart Tunnel Entry Dialog Boxes

Element	Description
App Name	The name of the application to which you are allowing smart tunnel access. The name can be up to 64 characters. Consider including the version number of the application if you are allowing more than one version smart tunnel access.

Table 34-41 Add and Edit Smart Tunnel Entry Dialog Boxes (continued)

Element	Description
App Path	The filename and optionally, the path, of the application. This entry can be up to 128 characters. Use one of the following: • Filename—For example, **outlook.exe**. By only specifying the file name, it does not matter where users install the application on their workstations. However, the file name must match exactly. • Full path and filename—For example, **C:\Program Files\Microsoft Office\OFFICE11\OUTLOOK.EXE**. This allows the application smart tunnel access only if it is installed in the specified directory, which you can use to enforce organizational standards. **Tips** • If you specify the full path, and the smart tunnel application stops working after it had been working for a while, it is likely that a product upgrade changed the installation path. Add a new entry that accounts for the new path. • If you are granting smart tunnel access to an application that is started from the command line, create one entry for **cmd.exe** (the Windows command line), and another entry for the application.
Platform	Specify the host operating system of the application: • Windows • Mac

Table 34-41 Add and Edit Smart Tunnel Entry Dialog Boxes (continued)

Element	Description
Hash Value	(Optional) The hash value for the application. By specifying a hash value, you can ensure that the user does not rename another application to use a supported filename and thus start an unsupported and undesired application over the smart tunnel. To obtain the hash value, enter the checksum of the application (that is, the checksum of the executable file) into a utility that calculates a hash using the SHA-1 algorithm. One example of such a utility is the Microsoft File Checksum Integrity Verifier (FCIV), which is available at http://support.microsoft.com/kb/841290/. Place a temporary copy of the application to be hashed on a path that contains no spaces (for example, c:\temp) and then enter **fciv.exe -sha1** application at the command line (for example, **fciv.exe -sha1 c:\msimn.exe**) to display the SHA-1 hash. Copy and paste the value into this field. The SHA-1 hash is always 40 hexadecimal characters. Before authorizing an application for smart tunnel access, clientless SSL VPN calculates the hash of the application matching the App Name. It qualifies the application for smart tunnel access if the result matches the value of hash. Because the checksum varies with each version or patch of an application, the hash you enter can match only one version or patch on the remote host. To specify a hash for more than one version of an application, create a unique smart tunnel entry for each hash value. **Tip** Hash values require maintenance. You must update the smart tunnel list if you want to support future versions or patches of an application for which you supply a hash value. A sudden problem with smart tunnel access might be an indication that the application list containing hash values is not up-to-date with an application upgrade. You can avoid this problem by not entering a hash.

Add and Edit Smart Tunnel Network Lists Dialog Boxes

Beginning from Security Manager version 4.7, you can use the Add and Edit Smart Tunnel Network Lists dialog boxes to create and edit a list of hosts that you can use for configuring smart tunnel policies.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Smart Tunnel Network Lists** from the Object Type selector. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**. Alternatively, you can click the **Add (+)**button to add a new object, or click the **Edit (pencil)** button to edit an object.

Related Topics

- ASA Group Policies SSL VPN Clientless Settings, page 34-12
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Policy Object Manager, page 6-4

- Add and Edit A Smart Tunnel Network List Entry Dialog Box, page 34-70

Field Reference

Table 34-42 Add and Edit Smart Tunnel Network Lists Dialog Boxes

Element	Description
Name	The smart tunnel network list object name that you use to apply to the tunnel policy. The name can be up to 64 characters. Spaces are not allowed. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the network list object.
Smart Tunnel Network List Entries table	The host mask or IP address of the network to which applications will be allowed smart tunnel access through the SSL VPN. • To add an entry, click the **Add Row** button to open the Add and Edit A Smart Tunnel Entry Dialog Boxes, page 34-67. • To edit an entry, select it and click the **Edit Row** button. • To delete an entry, select it and click the **Delete Row** button.
Include Other Lists	The other SSL VPN smart tunnel network list objects that you want to include in this object, if any. Enter the names of the objects or click **Select** to select them from a list or to create new objects. Separate multiple entries with commas.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit A Smart Tunnel Network List Entry Dialog Box

Use the Add and Edit A Smart Tunnel Network List Entry dialog box to create a new smart tunnel network list entry or edit an existing entry in the table in the SSL VPN Smart Tunnel Network Lists dialog box.

Navigation Path

From Add and Edit Smart Tunnel List Dialog Boxes, page 34-66, click the **Add Row** button beneath the Smart Tunnel Network List Entries table, or select an entry and click the **Edit Row** button.

Related Topics

- Add and Edit Smart Tunnel Network Lists Dialog Boxes, page 34-69
- ASA Group Policies SSL VPN Clientless Settings, page 34-12
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85

- Policy Object Manager, page 6-4

Field Reference

Table 34-43 Add and Edit Smart Tunnel Network List Entry Dialog Boxes

Element	Description
Host	The host mask that will be part of the smart tunnel network list entry.
IP Address	The IP address of the host that will be part of the smart tunnel network list entry. Beginning with version 4.12, Security Manager supports IPv6 addresses.
Subnet Mask	The subnet mask for the specified IP address.

Add and Edit Smart Tunnel Auto Signon List Dialog Boxes

Use the Add and Edit Smart Tunnel Auto Signon Lists dialog boxes to create, copy, and edit SSL VPN smart tunnel auto sign-on objects.

Smart Tunnel Auto Sign-on is a single sign-on method for Clientless SSL VPN users. It passes the login credentials (username and password) to internal servers for authentication using NTLM authentication, HTTP Basic authentication, or both. Smart Tunnel Auto Sign-on is supported on ASA 5500 devices running software version 7.1(1) and later.

An SSL VPN smart tunnel auto sign-on list object identifies the servers for which to automate the submission of login credentials during smart tunnel setup. You can configure the clientless settings of an ASA group policy with a smart tunnel auto sign-on list if you want to reissue the user credentials when the user establishes a smart tunnel connection to a server. For an explanation of the types of applications that support smart tunnel access, see Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85.

You can include other SSL VPN smart tunnel auto sign-on list objects in an object. Thus, you can create a set of objects that identify your basic list of servers and include those objects in another object that expands upon that list of servers.

Navigation Path

Select **Manage > Policy Objects**, then select **SSL VPN Smart Tunnel Auto Signon Lists** from the Object Type selector. Right-click inside the work area and select **New Object**, or right-click a row and select **Edit Object**.

Related Topics

- ASA Group Policies SSL VPN Clientless Settings, page 34-12
- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Policy Object Manager, page 6-4

Field Reference

Table 34-44 Add and Edit Smart Tunnel Auto Signon List Dialog Boxes

Element	Description
Name	The object name, which can be up to 64 characters. Spaces are not allowed. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Smart Tunnel Auto Signon Entries table	The servers for which to automate the submission of login credentials during smart tunnel setup. • To add servers, click the **Add Row** button to open the Add and Edit Smart Tunnel Auto Signon Entry Dialog Boxes, page 34-72. • To edit an entry, select it and click the **Edit Row** button. • To delete an entry, select it and click the **Delete Row** button.
Include Other Lists	The other smart tunnel auto sign-on list objects that you want to include in this object, if any. Enter the names of the objects or click **Select** to select them from a list or to create new objects. Separate multiple entries with commas.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add and Edit Smart Tunnel Auto Signon Entry Dialog Boxes

Use the Add and Edit Smart Tunnel Auto Signon Entry dialog boxes to create a new smart tunnel entry or edit an existing entry in the table in the SSL VPN Smart Tunnel Auto Signon List dialog box.

Navigation Path

From Add and Edit Smart Tunnel Auto Signon List Dialog Boxes, page 34-71, click the **Add Row** button beneath the Smart Tunnel Auto Signon Entries table, or select an entry and click the **Edit Row** button.

Related Topics

- Configuring SSL VPN Smart Tunnels for ASA Devices, page 31-85
- Policy Object Manager, page 6-4

Field Reference

Table 34-45 Add and Edit Smart Tunnel Auto Signon Entry Dialog Boxes

Element	Description
Matching Mode: • Host • IPv4/IPv6 Address	Identifies the server for which to automate the submission of login credentials during smart tunnel setup. Use Host to specify the server by host name or wildcard mask, and use IP Address to specify the server by IP address and netmask: • **Host**—Select Host and then enter the host name or a wildcard mask in the Hostname Mask field that identifies the host for which to automate the submission of login credentials during smart tunnel setup. **Note** Using this option protects the configuration from dynamic changes to IP addresses. • **IPv4/IPv6 Address**—Select the IP Address and then enter the IP address and netmask of the host or sub-network of hosts for which to automate the submission of login credentials during smart tunnel setup. **Note** Beginning with version 4.12 Security Manager supports IPv6 addresses. By default, when you select the **IPv4/IPv6 Address,** Security Manager looks for IPv4/IPv6 address. Enter the Subnet Mask or Prefix Length as required. **Note** Firefox requires the administrator to specify hosts using an exact host name or IP address (instead of a host mask with wild cards, a subnet using IP addresses, or a netmask). For example, within Firefox, you cannot enter *.cisco.com and expect auto sign-on to host email.cisco.com.
Port Number	The port that performs auto sign-on. For Firefox, if no port number is specified, auto sign-on is performed on HTTP and HTTPS, accessed by default port numbers 80 and 443 respectively.
Authentication Realm	The realm for the authentication. The Authentication Realm is associated with the protected area of the website and is passed back to the browser either in the authentication prompt or in the HTTP headers during authentication. After auto sign-on is configured and a realm string is specified, users can configure the realm string on a web application (such as Outlook Web Access) and access web applications without signing on.
Use Domain	Select this option to add the Windows domain to the username if authentication requires it. If you use this option, be sure to specify the domain name when assigning the smart tunnel list to one or more group policies.

Add or Edit User Group Dialog Box

Use the Add or Edit User Group dialog box to create or edit a user group object. User group objects are used in Easy VPN topologies, remote access VPNs, and SSL VPNs for IOS devices.

When you configure a remote access VPN, SSL VPN, or Easy VPN server, you can create user groups to which remote clients belong. The remote clients must be configured with the same group name as the user group on the VPN server in order to connect to the server; otherwise, no connection is established. When the remote client connects to the VPN server successfully, the group policies for that particular user group are pushed to all remote clients belonging to the user group.

For more information about user groups, see:

- Configuring User Group Policies, page 33-13
- Configuring a User Group Policy for Easy VPN, page 28-14
- Configuring an SSL VPN Policy (IOS), page 33-14

Note You must select the technology (Easy VPN/Remote Access VPN, or SSL VPN) for which you are creating the user group object. If you are editing an existing user group object, the technology is already selected and you cannot change it. Depending on the selected technology, the appropriate settings are available for configuration.

Navigation Path

Select **Manage > Policy Objects**, then select **User Groups** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Tip You can also access this dialog box from the **Remote Access VPN > IPSec VPN > User Groups** or the **Remote Access VPN > SSL VPN** policies.

Related Topics

- Policy Object Manager, page 6-4

Field Reference

Table 34-46 User Group Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Settings Pane

The body of the dialog box is a pane with a table of contents on the left and settings related to the item selected in the table of contents on the right.

You must first configure technology settings, then you can select items from the table of contents on the left and configure the options you require. Your selections on the Technology page control which options are available on these pages and in the table of contents.

The top folders in the table of contents represent the VPN technologies or other settings that you can configure, and are explained next.

Table 34-46 User Group Dialog Box (continued)

Element	Description
Technology settings	These settings control what you can define in the group policy: • **Group Name**—The name for the user group (up to 128 characters). Configure the same user group name within the remote client or device to ensure that the appropriate group attributes are downloaded. • **Technology**—The types of VPN for which this object defines group policies. You cannot change this option when editing an object, or if you are creating the user group object while editing a VPN policy. You can configure settings for Easy VPN/Remote Access IPSec VPN or SSL VPN, but not both.
Easy VPN/Remote Access IPSec VPN pages	When you select Easy VPN/Remote Access IPSec VPN as the technology, you can configure settings on the following pages: • User Group Dialog Box—General Settings, page 34-75 • User Group Dialog Box—DNS/WINS Settings, page 34-77 • User Group Dialog Box—Split Tunneling, page 34-77 • User Group Dialog Box—IOS Client Settings, page 34-78 • User Group Dialog Box—IOS Xauth Options, page 34-80 • User Group Dialog Box—IOS Client VPN Software Update, page 34-81 • User Group Dialog Box—Advanced PIX Options, page 34-82
SSL VPN pages	When you select SSL VPN as the technology, you can configure settings on the following pages: • User Group Dialog Box—Clientless Settings, page 34-83 • User Group Dialog Box—Thin Client Settings, page 34-84 • User Group Dialog Box—SSL VPN Full Tunnel Settings, page 34-84 • User Group Dialog Box—DNS/WINS Settings, page 34-77 • User Group Dialog Box—SSL VPN Split Tunneling, page 34-86 • User Group Dialog Box—Browser Proxy Settings, page 34-87 • User Group Dialog Box—SSL VPN Connection Settings, page 34-88
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

User Group Dialog Box—General Settings

The general settings you configure for your user group include the authentication method, IP address pool information, and connection attributes for PIX 6.3 Firewalls.

Note These settings apply in Easy VPN and remote access IPSec VPN configurations.

Navigation Path

Select **General** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-47 User Group Dialog Box—General Settings

Element	Description
Preshared Key	The preshared key that will be used to authenticate the clients associated to the user group. **Note** You do not have to enter a preshared key if you are using digital certificates for group authentication. In regular IPsec VPNs, preshared keys allow for one or more peers to use individual shared secrets to authenticate encrypted tunnels. A preshared key must be configured on each participating peer. If one of the participating peers is not configured with the same preshared key, the IKE SA cannot be established. In Easy VPN authentication, the same Easy VPN server key is used for the spoke configuration to ensure that the server/client keys match. In remote access IPSec VPN authentication, the same key is used to negotiate a VPN connection between the remote access VPN server and the remote clients.
IP Address Pool Subnet/Ranges	The IP address ranges for a local pool that will be used to allocate an internal IP address to a client. Remote clients are assigned IP addresses from this pool. Separate multiple entries with commas. The default is 172.16.0.1-172.16.4.254.
Backup Servers IP Address	The IP address of the servers to be used as backups for the Easy VPN or remote access IPSec VPN server. The router tries to connect to these servers if the primary connection to the Easy VPN or remote access VPN server fails. Separate multiple entries with commas.
PIX Only Attributes	These attributes apply only to PIX 6.3 devices. • **Idle Time**—The timeout period for VPN connections, in seconds. If no communication occurs on the connection during this period, the device terminates the connection. The minimum is 60 seconds, and the maximum time is 35791394 minutes. The default is 30 minutes. • **Max Time**—The maximum amount of time for VPN connections, in seconds. At the end of the time, the device terminates the connection. The minimum is 60 seconds, and the maximum is 35791394 minutes. There is no default.

User Group Dialog Box—DNS/WINS Settings

Configure the DNS/WINS settings for your user group to define the DNS and WINS servers and the domain name that should be pushed to clients associated with the user group.

Note The DNS/WINS settings you configure for a user group apply in Easy VPN, remote access VPN, and SSL VPN configurations.

Navigation Path

Select **DNS/WINS** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-48 User Group Dialog Box—DNS/WINS Settings

Element	Description
Primary DNS Server	The IP address of the primary DNS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary DNS Server	The IP address of the secondary DNS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Domain Name	The domain name of the DNS server you want to configure on the user group.
Primary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.
Secondary WINS Server	The IP address of the primary WINS server for the group. Enter the IP address or the name of a network/host object, or click **Select** to select an object from a list or to create a new object.

User Group Dialog Box—Split Tunneling

Split tunneling lets a remote client conditionally direct packets over an IPsec or SSL VPN tunnel in encrypted form or to a network interface in clear text form. With split tunneling enabled, packets not bound for destinations on the other side of the tunnel do not have to be encrypted, sent across the tunnel, decrypted, and then routed to a final destination.

The split tunneling policy is applied to a specific network. When you configure split tunneling, you can transmit both secured and unsecured traffic on the same interface. You must specify which traffic will be secured and what the destination of that traffic is, so that you have a secure tunnel to the central site, while the clear (unsecured) traffic is transmitted across the public network.

Tip For optimum security, we recommend that you not enable split tunneling.

Note Split tunneling can be applied in Easy VPN, remote access VPN, and SSL VPN configurations. For information about configuring split tunneling for SSL VPN, see User Group Dialog Box—SSL VPN Split Tunneling, page 34-86.

Navigation Path

Select **Split Tunneling** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73 when configuring Easy VPN/Remote Access IPSec VPN.

Field Reference

Table 34-49 User Group Dialog Box—Split Tunneling

Element	Description
Split Tunneling	The networks for which you want to tunnel traffic. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. You can identify the networks using one of these options: • **Protected Networks**—Specify the networks by network addresses. Enter the addresses or network/host objects, or click **Select** to select the objects from a list or to create new objects. For information on specifying addresses, see Specifying IP Addresses During Policy Definition, page 6-87. • **ACL**—Specify the networks using an extended access control list policy object. Enter the name of the object or click **Select** to select the object from a list or to create a new object.
Split DNS	A list of domain names that must be tunneled or resolved to the private network. All other names will be resolved through the public DNS server. You can enter multiple domain names separated by commas.

User Group Dialog Box—IOS Client Settings

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Configure IOS client settings to define Cisco IOS specific options for your user group, including firewall settings for VPN clients.

Note These settings apply in Easy VPN and remote access IPSec VPN configurations.

Navigation Path

Select **Client Settings (IOS)** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-50 User Group Dialog Box—Client Settings (IOS)

Element	Description
Enable Firewall Are-You-There (Not available on 7600 series or ASR routers.)	This feature may be used if a VPN client is running the Black Ice or Zone Alarm personal firewall. When selected, it ensures that the personal firewall is running at connection time and throughout the connection. The Firewall-Are-U-There attribute is sent by the Black Ice and Zone Alarm personal firewalls if the server prompts them to do so. If the personal firewall stops running, the connection is terminated. If this feature is enabled and there is no personal firewall running on the server, the connection is never established.
Mode	A Central Policy Push (CPP) firewall policy on a server allows or denies a tunnel on the basis of whether the remote device has a required firewall for a local AAA server. The Mode option specifies whether the Central Policy Push (CPP) policy is optional or mandatory, as follows: • **Optional**—If the CPP policy is defined as optional, and is included in the Easy VPN server configuration, the tunnel setup is continued even if the client does not confirm the defined policy. • **Required**—If the CPP policy is defined as mandatory and is included in the Easy VPN server configuration, the tunnel setup is allowed only if the client confirms this policy. Otherwise, the tunnel is terminated.
Firewall Type	The type of firewall that you are making required or optional. The list shows all of the supported firewall software, which includes software from Cisco and Zone Labs.
Policy Type	Specifies the CPP firewall policy type: • **Check Presence**—Instructs the server to check for the presence of the specified firewall type. • **Central Policy Push**—The actual policy, such as the input and output access lists, that must be applied by the specified client firewall type. Specify the following: – The access control list to be used. Enter the name of the extended ACL object or click **Select** to select it from a list or to create a new object. – The direction of the access control list—Inbound or Outbound.
Include Local LAN	Whether to allow a non split-tunneling connection to access the local LAN at the same time as the client.
Perfect Forward Secrecy	Whether to enable Perfect Forward Secrecy (PFS). If PFS is enabled, the server is configured to notify the client of the central-site policy about whether PFS is required for any IPsec SA. The Diffie-Hellman (D-H) group that is proposed for PFS is the same that was negotiated in Phase 1 of the IKE negotiation.

User Group Dialog Box—IOS Xauth Options

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

IOS Xauth options configure IKE Extended Authentication (Xauth) user authentication and connection parameters for the user group, including the banner text.

Note These settings apply in Easy VPN and remote access VPN configurations.

Navigation Path

Select **Xauth Options (IOS)** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-51 User Group Dialog Box—IOS Xauth Options

Element	Description
Banner	The banner text that is displayed to Easy VPN remote clients during Xauth and web-based activation the first time the Easy VPN tunnel is brought up. A maximum of 1024 characters is allowed.
Maximum Logins Per User	The maximum number of connections a user can establish simultaneously. The maximum is 10.
Maximum Connections	The maximum number of client connections to the Easy VPN Server from this group. The maximum is 5000 per group.
Enable Group-Lock	Whether to enable group lock, which requires that the user enter the extended Xauth username in one of the following formats: • username/groupname • username\groupname • username@groupname • username%groupname The group that is specified after the delimiter is then compared to the group identifier that is sent during IKE aggressive mode. The groups must match or the connection is rejected. **Note** Do not select this option if you are using RSA signature authentication mechanisms such as certificates.
Enable Save Password	Whether to allow users to save their Xauth password locally on the client. On subsequent authentications, users can activate the password by using the check box on the software client or by adding the username and password to the Cisco IOS hardware client profile. After users activate the password, their username and password are sent to the server automatically during Xauth. This option is useful only if users have static passwords, that is, they are not one-time passwords such as those that are generated by a token.

User Group Dialog Box—IOS Client VPN Software Update

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any enhancements.

Client VPN Software Update (IOS) settings configure, for an IOS VPN client, the platform type, VPN Client revisions, and image URL for each client VPN software package installed, for your user group.

The Client Update feature is supported on IOS routers version 12.4(2)T and later, and Catalyst 6500/7600 devices version 12.2(33)SRA and later.

- To add a client, click the **Add Row** button to open the Add/Edit Client Update Dialog Box, page 34-81.
- To edit a client, select it and click the **Edit Row** button.
- To delete a client, select it and click the **Delete Row** button.

Note These settings apply in Easy VPN and remote access VPN configurations.

Navigation Path

Select **Client VPN Software Update (IOS)** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Add/Edit Client Update Dialog Box

Use the Add or Edit Client Update dialog box to configure the platform type, image URL, and VPN Client revisions for a client VPN software package.

Navigation Path

Open the User Group Dialog Box—IOS Client VPN Software Update, page 34-81, then click **Add Row**, or select an item in the table and click **Edit Row**.

Related Topics

- Add or Edit User Group Dialog Box, page 34-73

Field Reference

Table 34-52 Add or Edit Client Update Dialog Box

Element	Description
System Type	The platform on which the IOS VPN client operates. • All Windows (Default)—This option includes any Windows platform for which a VPN client is available. • Macintosh OS X
IOS Image URL	Enter the URL from where the client can be downloaded. The URL must start with http:// or https://.
IOS VPN Client Revisions	Enter the revision level of the VPN client. You can specify more than one client revision separated by commas.

User Group Dialog Box—Advanced PIX Options

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any enhancements.

The Advanced PIX Options are specifically for PIX 6.3 Firewalls in your user group.

Note These settings apply in Easy VPN and remote access VPN configurations.

Navigation Path

Select **Advanced Options (PIX)** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-53 User Group Dialog Box—Advanced PIX Options

Element	Description
User Idle Timeout (sec)	The length of time that a VPN tunnel can remain open without user activity, in seconds. Values range from 60-86400 seconds.
User Authentication Server	The AAA server to which remote devices send user authentication requests. Enter the name of the server group or click **Select** to select it from a list or to create a new group. See Understanding AAA Server and Server Group Objects, page 6-27.
Enable Device Pass-Through	Whether to use Media Access Control (MAC) addresses to bypass authentication for devices, such as Cisco IP phones, that do not support AAA authentication. When MAC-based AAA exemption is enabled, the device bypasses the AAA server for traffic that matches both the MAC address of the device and the IP address that was dynamically assigned by a DHCP server. Authorization services are disabled automatically when you bypass authentication. Accounting records continue to be generated (if enabled), but the username is not displayed.
Enable Secure Unit Authentication	Whether to provide increased security when allowing access to the device from a remote client. With Secure Unit Authentication (SUA), you can use one-time passwords, two-factor authentication, and similar authentication schemes to authenticate the remote device during Extended Authentication (Xauth). SUA is specified in the VPN policy on the device and is downloaded to the remote client. This enables SUA and determines the connection behavior of the remote client.
Enable User Authentication	Whether to enable Individual User Authentication (IUA), which supports individually authenticating clients on the inside network of the remote access VPN, based on the IP address of each inside client. IUA supports both static and OTP authentication mechanisms.

User Group Dialog Box—Clientless Settings

Use the Clientless settings to configure the clientless mode of access to the corporate network in an SSL VPN.

In clientless access mode, once a user is authenticated and a session is established, an SSL VPN portal page and toolbar is displayed on the user's web browser. From the portal page, the user can access all available HTTP sites, access web e-mail, and browse Common Internet File System (CIFS) file servers.

Navigation Path

Select **Clientless** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Related Topics

- Create Group Policy Wizard—Clientless and Thin Client Access Modes Page, page 30-24

Field Reference

Table 34-54 User Group Dialog Box—Clientless Settings

Element	Description
Portal Page Websites	The name of the SSL VPN bookmarks policy object that includes the web site URLs to display on the portal page. These web sites help users access desired resources. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Allow Users to Enter Websites	Whether to allow the remote user to enter web site URLs directly into the browser. If you do not select this option, the user can access only those URLs included on the portal.
Enable Common Internet File System (CIFS)	In Clientless mode, files and directories created on Microsoft Windows servers can be accessed by the remote client through the web browser. When you enable the Common Internet File System (CIFS), a list of file server and directory links are displayed on the portal page after login. The CIFS protocol lets you customize permissions on the SSL VPN gateway to allow shared files to be accessed or modified by the remote client, as follows: • **Enable File Browsing**—Whether to allow the remote user to browse for file shares on the CIFS file servers. • **Enable File Entry**—Whether to allow the remote user to locate file shares on the CIFS file servers by entering the names of the file shares.
WINS Server List	The name of the WINS server list policy object that identifies the WINS/NetBIOS servers to use for resolving file server names. You should supply an object if you enable CIFS. Enter the name of the object or click **Select** to select if from a list or to create a new object.
Enable Citrix	Whether to enable remote clients to run Citrix-enabled applications, such as Microsoft Word or Excel, through the SSL VPN as if the application were locally installed, without the need for client software. The Citrix software must be installed on one or more servers on a network that the router can reach.

User Group Dialog Box—Thin Client Settings

Use the Thin Client settings to enable the thin client, or port forwarding, mode of access to the corporate network in an SSL VPN. Port forwarding allows users to access applications (such as Telnet, e-mail, VNC, SSH, and Terminal services) inside the enterprise through an SSL VPN session. A port forwarding list object defines the mappings of port numbers on the remote client to the application's IP address and port behind the SSL VPN gateway.

In thin client access mode, the remote user downloads a Java applet that acts as a TCP proxy on the client machine for the services configured on the SSL VPN gateway. The proxy provides the port forwarding services.

Navigation Path

Select **Thin Client** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Related Topics

- Create Group Policy Wizard—Clientless and Thin Client Access Modes Page, page 30-24

Field Reference

Table 34-55 User Group Dialog Box—Thin Client Settings

Element	Description
Enable Thin Client	Whether to allow thin client access to the SSL VPN.
Port Forward List	The name of the port forwarding list policy object assigned to this group. Port forwarding lists contain the set of applications that users of clientless SSL VPN sessions can access over forwarded TCP ports. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Download Port Forwarding Applet on Client Login	Whether the port forwarding Java applet should be automatically downloaded to the client when a user logs into the SSL VPN. If you do not automatically download the applet, users must download it manually after login.

User Group Dialog Box—SSL VPN Full Tunnel Settings

Use the SSL VPN Full Tunnel settings to enable the full tunnel client access mode in your SSL VPN. When you enable full tunnel access, you should also define DNS/WINS server settings, browser proxy settings, and split tunneling for the user group.

In full tunnel client access mode, the tunnel connection is determined by the group policy configuration. The full tunnel client software, SSL VPN Client (SVC), must be downloaded to the remote client so that a tunnel connection can be established when the remote user logs in to the SSL VPN gateway.

Tip For full tunnel client access to work, you must install the client software on the gateway. The user downloads the client when connecting to the gateway.

Navigation Path

Select **Full Tunnel > Settings** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Related Topics

- Create Group Policy Wizard—Full Tunnel Page, page 30-21

Field Reference

Table 34-56 User Group Dialog Box—Full Tunnel Settings

Element	Description
Enable Full Tunnel	Whether to enable full tunnel client access to the SSL VPN.
Use Other Access Modes if SSL VPN Client Download Fails Full Tunnel Only	Whether to allow users to connect to the SSL VPN even if a problem prevents the client from downloading, installing, and starting correctly on the user's system. If you select **Full Tunnel Only**, a user cannot connect to the SSL VPN if the download fails, which locks the user out of the network. Select **Use Other Access Modes** to allow clientless or thin client access if there is a download problem.
Client IP Address Pool	The IP address ranges of the address pool that full tunnel clients will draw from when they log on. The address pool must be in the same subnet as one of the device's interface IP addresses. Enter the address range separating the first and last IP address with a hyphen, for example, **10.100.10.2-10.100.10.255**. If you enter a single address, the pool has just one address. Do not enter subnet designations. You can also enter the name of a network/host policy object that defines the range, or click **Select** to select the object from a list or to create a new object. Separate multiple ranges with commas.
Filter ACL	The name of an extended access control list (ACL) object that restricts access to the SSL VPN. Enter the name of the object or click **Select** to select it from a list or to create a new object.
Keep SSL VPN Client on Client Computer	Whether to leave the full client installed on the user's workstation after the user disconnects. If you do not allow the client to remain on the user's system, the client must be downloaded each time the user establishes a connection to the SSL VPN gateway.
Home Page URL	The web address of the login home page for the full client.
Client Dead Peer Detection Timeout	The time interval that the Dead Peer Detection (DPD) timer is reset each time a packet is received over the SSL VPN tunnel from the remote user. Enter a value in the range 1-3600 seconds.
Gateway Dead Peer Detection Timeout	The time interval that the Dead Peer Detection (DPD) timer is reset each time a packet is received over the SSL VPN tunnel from the gateway. Enter a value in the range 1-3600 seconds.

Table 34-56 User Group Dialog Box—Full Tunnel Settings (continued)

Element	Description
Key Renegotiation Method	The method by which the tunnel key is refreshed for the remote user group client: • **Disabled**—Disables the tunnel key refresh. • **Create New Tunnel**—Initiates a new tunnel connection. Enter the time interval (in seconds) between the tunnel refresh cycles in the **Interval** field.

User Group Dialog Box—SSL VPN Split Tunneling

Use the Split Tunneling settings to configure a secure tunnel to the central site and simultaneous clear text tunnels to the Internet for SSL VPNs.

Split tunneling lets a remote client conditionally direct packets over an IPsec or SSL VPN tunnel in encrypted form or to a network interface in clear text form. With split tunneling enabled, packets not bound for destinations on the other side of the tunnel do not have to be encrypted, sent across the tunnel, decrypted, and then routed to a final destination. The split tunneling policy is applied to specific networks.

Tip For optimum security, we recommend that you not enable split tunneling.

Navigation Path

Select **Full Tunnel > Split Tunneling** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-57 User Group Dialog Box—Split Tunneling Settings

Element	Description
Tunnel Option	Whether to allow split tunneling and if so, which traffic should be secured or transmitted unencrypted across the public network: • Disabled—(Default) No traffic goes in the clear or to any other destination than the gateway. Remote users reach networks through the corporate network and do not have access to local networks. • Tunnel Specified Traffic—Tunnel all traffic from or to the addresses listed in the **Destinations** field. Traffic to all other addresses travels in the clear and is routed by the remote user's Internet service provider. • Exclude Specified Traffic—Traffic goes in the clear from and to the addresses listed in the **Destinations** field. This is useful for remote users who want to access devices on their local network, such as printers, while they are connected to the corporate network through a tunnel.

Table 34-57 User Group Dialog Box—Split Tunneling Settings (continued)

Element	Description
Destinations	The IP addresses for hosts or networks that identify the networks that require traffic to travel across the tunnel and those that do not require tunneling. Whether traffic to these addresses is encrypted and tunneled to the gateway, or sent in the clear, is determined by your selection for **Tunnel Option**. Enter network addresses such as 10.100.10.0/24 or host addresses such as 10.100.10.12. You can also enter the name of a network/host policy object, or click **Select** to select the object from a list or to create a new object. Separate multiple addresses with commas.
Exclude Local LANs	Whether to exclude local LANs from the encrypted tunnel. This option is available only if you selected the **Exclude Specified Traffic** tunnel option. By selecting this option, you do not have to enter local LAN addresses into the destinations field to allow users to communicate with systems (such as printers) that are attached to their LAN. When selected, this attribute disallows a non split-tunneling connection to access the local subnetwork at the same time as the client.
Split DNS Names	A list of domain names to be resolved through the split tunnel to the private network. All other names are resolved using the public DNS server. Enter up to 10 entries in the list of domains, separated by commas. The entire string can be no longer than 255 characters.

User Group Dialog Box—Browser Proxy Settings

Use the Browser Proxy settings to configure proxy bypass for full tunnel access in an SSL VPN.

A security appliance can terminate HTTPS connections and forward HTTP/HTTPS requests to HTTP and HTTPS proxy servers, which act as intermediaries between users and the Internet. Proxy bypass is an alternative method of content rewriting that makes minimal changes to the original content. It is useful with custom web applications.

Tip The browser proxy settings work only for Microsoft Internet Explorer; they do not work for other types of browsers.

Navigation Path

Select **Full Tunnel > Browser Proxy Settings** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Related Topics

- Configuring SSL VPN Proxies and Proxy Bypass (ASA), page 31-57

Field Reference

Table 34-58 User Group Dialog Box—Browser Proxy Settings

Element	Description
Browser Proxy Option	Whether and how to configure proxy settings on the remote client's browser: • Blank—Do not configure proxy settings. • Do Not Use Proxy Server—Configure the browser to not use a proxy. • Automatically Detect Settings—Configure the browser to automatically detect proxy settings. • Bypass Proxy Server for Local Addresses—Configure the browser to bypass proxy settings configured by the user.
Proxy Server	The address of the proxy server: • IP address—The IP address or the name of a network/host object that specifies the address. Click **Select** to select the object from a list. • Name—The fully qualified domain name, for example, proxy.example.com.
Proxy Server Port	The port number on the server that is used for proxy traffic, for example, 80. Enter a value in the range 1-65535.
Do Not Use Proxy Server for Addresses Beginning With	If you configured a proxy, you can identify specific hosts for which the proxy should be bypassed. If the user opens these hosts in the browser, the proxy is not used in the connection. Enter full IP addresses or fully qualified domain names. For example, 10.100.10.14 or www.cisco.com.

User Group Dialog Box—SSL VPN Connection Settings

Use this SSL VPN Connection Settings page to configure the SSL VPN session connection settings for the user group, including the banner text. An SSL VPN session is disconnected if the client is connected longer than the session timeout or if it is idle longer than the idle timeout.

Navigation Path

Select **Connection Settings** from the table of contents in the Add or Edit User Group Dialog Box, page 34-73.

Field Reference

Table 34-59 User Group Dialog Box—Connection Settings

Element	Description
Idle Timeout	The idle timeout period for the SSL VPN session. The session is disconnected if the client is idle longer than the specified idle timeout. Values range from 0-3600 seconds.

Table 34-59 User Group Dialog Box—Connection Settings (continued)

Element	Description
Session Timeout	The timeout period for the SSL VPN session. The session is disconnected when this timeout is reached even if the user is still active. Values range from 1-1209600 seconds.
Banner Text	The banner, for example, a welcome message, that is displayed to remote users when they connect to the SSL VPN. You cannot use double quotes or new lines (carriage returns) in the banner text. However, you can include HTML tags to create the desired layout.

Add or Edit WINS Server List Dialog Box

Use the WINS Server Lists dialog box to create, copy, and edit WINS server list objects. A WINS Server List object defines a list of Windows Internet Naming Server (WINS) servers, which are used to translate Windows file server names to IP addresses.

Navigation Path

Select **Manage > Policy Objects**, then select **WINS Server Lists** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs, page 31-88
- Policy Object Manager, page 6-4

Field Reference

Table 34-60 WINS Server Lists Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
WINS Server List	The WINS servers that are defined for the object. • To add a server, click the Add button and fill in the Add WINS Server dialog box (see Add or Edit WINS Server Dialog Box, page 34-90). • To edit a server, select it and click the Edit button. • To delete a server, select it and click the Delete button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 34-60 WINS Server Lists Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit WINS Server Dialog Box

Use the Add/Edit WINS Server dialog box to create a new WINS server entry or edit an existing entry in the table in the WINS Server Lists dialog box.

Navigation Path

From the Add or Edit WINS Server List Dialog Box, click the **Add** button beneath the WINS Server List table, or select a server in the table and click the **Edit** button.

Related Topics

- Configuring WINS/NetBIOS Name Service (NBNS) Servers To Enable File System Access in SSL VPNs, page 31-88

Field Reference

Table 34-61 Add/Edit WINS Server Dialog Box

Element	Description
Server	The IP address of the WINS server used to translate Windows file server names to IP addresses. You can also enter the name of a network/host policy object that identifies the server. Click **Select** to choose a network/hosts object or to create a new object.
Set as Master Browser	Whether to set the server as a primary browser. The primary browser maintains the list of computers and shared resources.
Timeout	The period of time the security appliance waits for a response to a WINS query before sending the query again to the same server (if it is the only one), or to the next server (if there is more than one). The default timeout is 2 seconds. The range is between 1 and 30 seconds.
Retries	The number of times to retry sending WINS queries to the configured servers. The security appliance recycles through the list of servers this number of times before sending an error message. The default is 2. The range is between 0 and 10.

Using Map View

The following topics describe how to use the Map view:

- Understanding Maps and Map View, page 35-1
- Working With Maps, page 35-8
- Displaying Your Network on the Map, page 35-14
- Managing VPNs in Map View, page 35-20
- Managing Device Policies in Map View, page 35-22

Understanding Maps and Map View

The Security Manager Map view provides a graphical view of your VPN and Layer 3 network topology.

Using the map view, you can investigate details of your VPN configuration graphically. Topological display of tunnels enables you to easily derive the relationship among multiple VPN configurations (for example, a hierarchical VPN). You can group devices to achieve a more complete picture of your VPN configuration. This is useful in situations where a hub failover pair is a peer with hundreds of spokes.

You can represent your Layer 3 network topology graphically, populating it with managed devices (called device nodes). You can make the picture of the topology more complete by adding unmanaged objects (called map objects) such as devices, clouds, and networks. For large networks, you can choose to simplify the topology graph by incorporating only a portion of the overall topology. You can save the topology maps for future use.

You can save multiple topology maps to reflect your network's geographical or functional organization. You can link a saved map to a node on a parent map, so that from the parent map you can drill down to the linked map with more detailed information (for more information, see Using Linked Maps, page 35-13). Saved maps are shared among all users who have the necessary access privileges.

You can launch other Security Manager features from the map view. In some cases, you can simplify the use of features by selecting nodes from the map before you start another feature. For example, you can select multiple nodes, then create a VPN that includes those nodes as members.

Tip The network data that is displayed on maps is typically updated as this data changes. However, to be certain that a map displays current network data, you can refresh it manually by selecting **Map > Refresh Map**.

This section contains the following topics:

- Understanding the Map View Main Page, page 35-2
- Map Toolbar, page 35-4
- Using the Navigation Window, page 35-4
- Maps Context Menus, page 35-5
- Access Permissions for Maps, page 35-8

Understanding the Map View Main Page

Map view enables you to create customized, visual topology maps of your network, within which you can view connections between your devices and easily configure VPNs and access control settings. The following figure identifies the functional areas of the Map view.

To open the Map view main page, click the **Map View** button in the toolbar.

You can undock the map window, which enables you to use other product features while keeping the map open. To undock the map, select **Map > Undock Map View**. To dock the map window, select **Map > Dock Map View**.

Figure 35-1 Map View Main Page

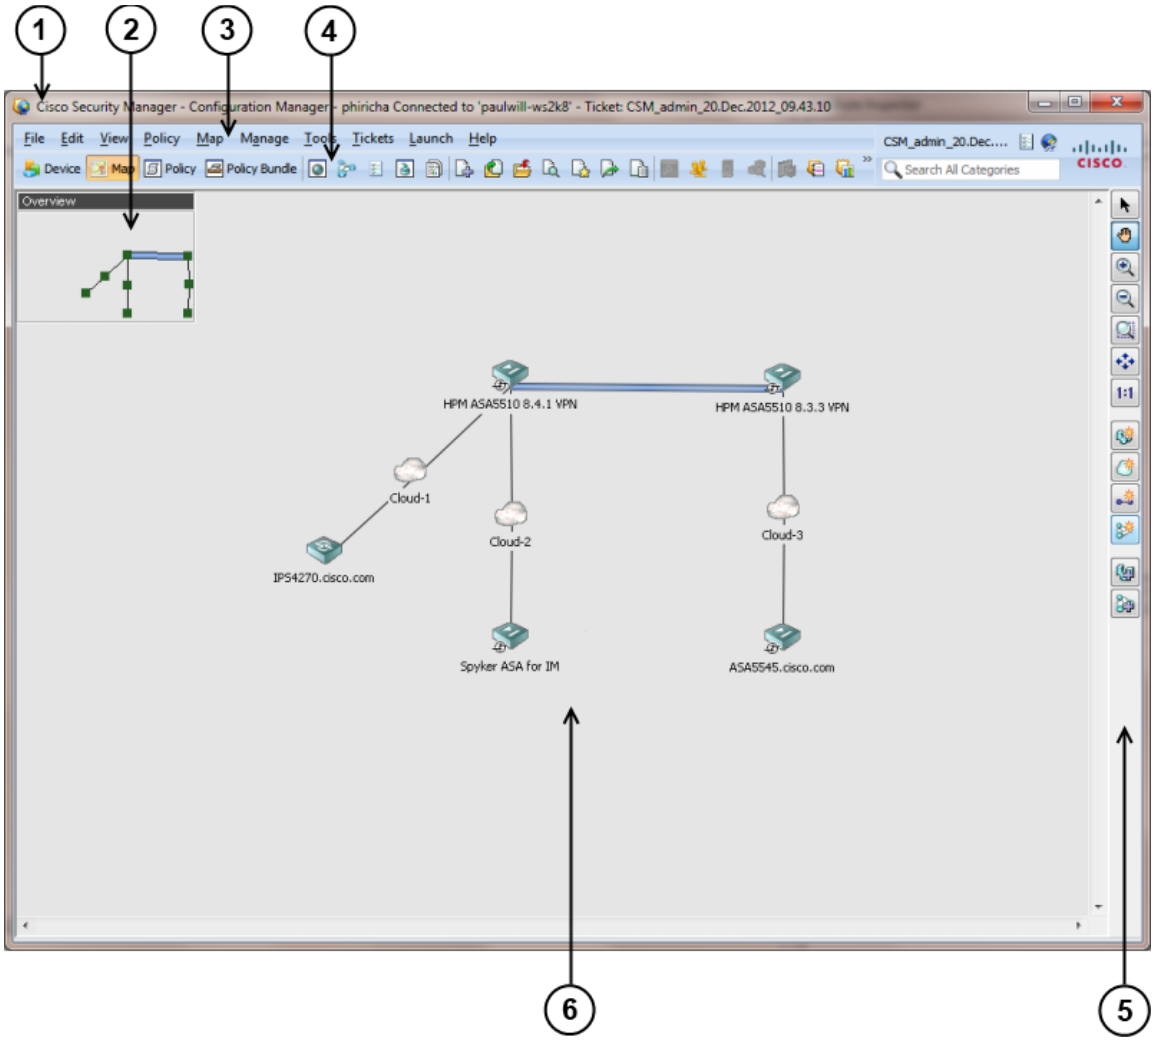

1	Title bar	**2**	Navigation window (see Using the Navigation Window, page 35-4)
3	Menu bar (see Map Menu (Configuration Manager), page 1-33)	**4**	Toolbar (see Toolbar Reference (Configuration Manager), page 1-39)
5	Map toolbar (see Map Toolbar, page 35-4)	**6**	Map (see Understanding Map Elements, page 35-14)

Related Topics

- Understanding Maps and Map View, page 35-1
- Working With Maps, page 35-8
- Displaying Your Network on the Map, page 35-14
- Managing VPNs in Map View, page 35-20
- Managing Device Policies in Map View, page 35-22

Map Toolbar

The following table describes the buttons on the map toolbar.

Table 35-1 Map Toolbar

Toolbar Button	Description
	Selects objects on the map. Click the button, then click items on the map.
	Pans the map. Click the button, click and hold on the map, then drag the cursor.
	Zooms in on the map.
	Zooms out from the map.
	Zooms the map to fill a rectangle that you draw.
	Zooms the map to include the entire map.
	Zooms the map to actual size.
	Creates a new Security Manager-managed node. After you create the new device in the inventory, it is added to the active map as a device node.
	Adds a new map object to the map.
	Adds a new link to the map.
	Creates a new VPN connection between nodes on the map.
	Select devices to show on the map as device nodes.
	Select VPNs to show on the map.

Using the Navigation Window

The navigation window displays a smaller version of the entire active map. The shaded rectangle defines the area of the map that is currently displayed.

Use the navigation window to select the portion of the map to view and to change the map zoom level.

- To toggle the display of the navigation window, select **Map > Show/Hide Navigation Window**.
- To pan the navigation control to select which portion of the map to display, click within the shaded rectangle and drag it to a new location.
- To change the zoom level, click one of the resizing handles in the corners of the shaded rectangle, then drag it to increase or decrease the area of the map to display. The map zooms to display the area covered by the map indicator.

The title bar in the navigation window displays the name of the map. If the map has unsaved changes, an asterisk (*) appears next to the map name.

For information on other ways to pan and zoom maps, see Panning, Centering, and Zooming Maps, page 35-11.

Maps Context Menus

The following topics describe the menus that contain maps commands. To open the context menus, right-click map elements.

- Managed Device Node Context Menu, page 35-5
- Multiple Selected Nodes Context Menu, page 35-6
- VPN Connection Context Menu, page 35-6
- Layer 3 Link Context Menu, page 35-7
- Map Object Context Menu, page 35-7
- Map Background Context Menu, page 35-7

Managed Device Node Context Menu

The Managed Device Node context menu opens when you right-click a map node that represents a managed device. The commands that you see depend on the type of device you select. The following table lists all commands that you might see.

Table 35-2 Managed Device Node Context Menu

Menu Command	Description
Edit Firewall Policies	Edits firewall policies on the device. Select a firewall policy type from the submenu to edit it.
Edit Firewall Settings	Edits firewall settings on the device. Select a setting from the submenu to edit it.
Edit VPN Peers	Edits peers in VPNs in which the device participates.
Edit VPN Policies	Edits VPN policies on the device.
Device Properties	Displays device properties.
Clone Device	Creates a copy of the device. See Cloning a Device, page 3-56.
Copy Policies Between Devices	Copies policies between the device and other devices. See Copying Policies Between Devices, page 5-33.
Share Device Policies	Shares device local policies.
Catalyst Summary Info	Allows you to view high-level system information, including any service modules, ports, and VLANs that Security Manager has discovered. See Viewing Catalyst Summary Information, page 68-2.
Show in Device View	Opens the Device View for the selected device.
Device Manager	Launches the Device Manager. See Starting Device Managers, page 72-14.

Table 35-2 Managed Device Node Context Menu (continued)

Menu Command	Description
Inventory Status	Displays the Inventory Status window for the device. See Inventory Status Window, page 72-13.
Show VPN Peers	Shows peers in VPNs in which the device participates.
Preview Configuration	Previews the device configuration with all committed changes included.
Show Containment	Shows the security contexts and service modules in devices that have them.
Node Properties	Displays node properties.
Set Linked Map	Creates a link from this node to another map.
Open Linked Map	Opens the map that is linked to the node.
Discover Policies on Device	Discovers policies on the device.
Move To Center	Pans the map to display the node in the center.
Delete Device	Deletes the device from the device inventory.
Remove from Map	Removes the node from the map.

Multiple Selected Nodes Context Menu

The Multiple Selected Device Node context menu opens when you select more than one map node, then right-click on a selected node.

If all of the selected nodes are not VPN-capable, the commands to configure VPNs do not appear.

Table 35-3 Multiple Selected Nodes Context Menu

Menu Command	Description
Create Point to Point VPN	Creates a point to point VPN between two selected devices. All selected nodes must be managed and VPN-capable.
Create Hub and Spoke VPN	Creates a hub and spoke VPN that includes the selected nodes. The node that you right-click becomes the VPN hub. All selected nodes must be managed and VPN-capable.
Create Meshed VPN	Creates a full mesh VPN that includes the selected nodes. All selected nodes must be managed and VPN-capable.
Remove Selected Nodes	Removes all selected device nodes. Appears only if you right-click on a selected device node.
Delete Map Objects	Deletes all selected map objects. Appears only if you right-click on a selected map object.

VPN Connection Context Menu

The VPN Connection context menu opens when you right-click on a VPN connection on the map. For more information, see Editing VPN Policies or Peers From the Map, page 35-22.

Table 35-4 VPN Connection Context Menu

Menu Command	Description
Edit VPN Peers	Edits the peers in the VPN.
Edit VPN Policies	Edits the VPN policies.

Layer 3 Link Context Menu

The Layer 3 Link context menu opens when you right-click on a layer 3 link on the map.

Table 35-5 Layer 3 Link Context Menu

Menu Command	Description
Link Properties	Displays the link properties.
Delete Link	Deletes the link from the map.

Map Object Context Menu

The Map Object context menu opens when you right-click a map object that does not represent a managed device.

Table 35-6 Map Object Context Menu

Menu Command	Description
Node Properties	Displays the node properties.
Move To Center	Pans the map to display the node in the center.
Set Linked Map	Links the node to a map.
Open Linked Map	Opens the map to which the node is linked.
Delete Map Object	Deletes the map object.

Map Background Context Menu

The Map Background context menu opens when you right-click in the background area of a map, that is, not on any object or link.

Table 35-7 Map Background Context Menu

Menu Command	Description
Show Devices on Map	Selects the managed devices to show on the map.
Show VPNs on Map	Selects the VPNs to display on the map.
Add Map Object	Adds a map object to the map.
Add Link	Adds a Layer 3 link to the map.
New Device	Creates a new managed device and adds it to the map as a device node.
New VPN	Creates a new VPN and adds it to the map.
Find Map Node	Finds nodes on the map.

Table 35-7 Map Background Context Menu (continued)

Menu Command	Description
Open Map	Opens a saved map.
Save Map	Saves the open map.
Show/Hide Navigation Window	Toggles the display of the navigation window on the map.
Map Properties	Displays the properties of the map.
Hierarchical layout	Arranges the network nodes in a hierarchical layout.
Radial layout	Arranges the network nodes in a radial layout.
Circular layout	Arranges the network nodes in a circular layout.
Dock/Undock Map	Undocks the Map view.

Access Permissions for Maps

Access to maps is controlled based on two systems of user privileges:

- Device privileges—You must have at least read privileges to all the devices in a map to open the map.
- Map privileges—Access to maps is based on your Security Manager user role. There are two levels of map access:
  - Read-only—You can open maps, but you cannot modify them. If you have this map privilege level, the features for modifying maps are not available.
  - Read-write—You can modify maps. All map modification features are available.

Working With Maps

A map is a representation of a portion of your network. You can create and save multiple maps to address your network management needs. To work with any map, you must be in Map view (select **View > Map View**).

After you create and save a map, the map is available to all users on the system that have at least read privileges to all the devices on the map. Users that do not have read privileges to a device on a map do not see the map in the list of existing maps when they try to open a map. For more information, see Access Permissions for Maps, page 35-8.

You can only have one map open at a time. If a map is open and you create a new map or open an existing map, you are prompted to save or discard any unsaved changes that you made to the current map.

Multiple users can open and modify a map at the same time. When a user saves changes to a map, any other users who are using the map are notified and have the option to do one of the following:

- Update their map to the version saved by the other user, losing any changes they have made.
- Save their version of the map as a new map, preserving any changes they made.

This section contains the following topics:

- Creating New or Default Maps, page 35-9
- Opening Maps, page 35-10

- Saving Maps, page 35-10
- Deleting Maps, page 35-10
- Exporting Maps, page 35-11
- Arranging Map Elements, page 35-11
- Panning, Centering, and Zooming Maps, page 35-11
- Selecting Map Elements, page 35-12
- Searching for Map Nodes, page 35-12
- Using Linked Maps, page 35-13
- Setting the Map Background Properties, page 35-13

Creating New or Default Maps

You have two options for creating a new map:

- Create an empty map—To create a new empty map, select **Map > New Map**. You must already be in Map view (select **View > Map View**). If you currently have a map open with unsaved changes, you are asked if you want to save it. For information about adding elements to a map, see Displaying Your Network on the Map, page 35-14.
- Create a new map containing all managed devices and VPNs in the inventory—This is called the default map. Generating the default map is a good way to create a map. After generating the map, save it with a unique name to make it a standard map and modify it as desired.

  You can generate the default map whenever you want to, and it contains the inventory as it exists at the time you generate it. You cannot specifically save the default map *as* the default map; it is regenerated every time you select it.

The following procedure explains how to create a new map using the default map.

Tips

- If you refresh the map (select **Map > Refresh Map**), items that you added to the inventory after generating the default map are not added to the map. You must reopen the default map to see new devices.

Step 1 In Map view, select **Map > Open Map**.

Step 2 Select **Default Map** from the Available Maps list, then click **OK**.

Note If you have do not have sufficient access rights to all devices in the inventory, the default map that opens shows only the subset of devices for which you do have access rights. For more information, see Access Permissions for Maps, page 35-8.

Step 3 To save the default map as a standard map, select **Map > Save Map** or **Map > Save Map As**, enter a name for the map and click **OK**.

Opening Maps

To open an existing map, select **Map > Open Map**, select the desired map from the list of available maps, and click **OK**. You must already be in Map view (select **View > Map View**). If you currently have a map open with unsaved changes, you are asked if you want to save it.

The list of available maps includes a special map called the **Default Map**. This map contains all of the managed devices and VPNs in the inventory. You are essentially creating a new map each time you open it. For more information about the default map, see Creating New or Default Maps, page 35-9.

Tip You can open any map that you have created or the default map. You can also open any map that another user has created provided you have the requisite permission settings with regard to the devices shown on that map (see Access Permissions for Maps, page 35-8).

Related Topics

- Working With Maps, page 35-8
- Understanding Map Elements, page 35-14

Saving Maps

To save the active map, select **Map > Save Map**. Any changes that you made since you last saved it are saved. If you did not save the map previously, the Save Map As dialog box opens, enabling you to assign a name to the map and save it.

To save a map under a new name, select **Map > Save Map As**. The map name can be as long as 256 characters, but cannot be the reserved names “Default Map” or “New Map.”

If you close a map that contains unsaved changes, you are prompted to save the changes.

If your Security Manager session closes automatically because of inactivity when a map is open with unsaved changes, the current version of the map is saved if it has a name. If you have not yet saved the map, the map is discarded. For example, if you generate the default map, or create a new map, and do not save it before your session times out, you cannot retrieve that map.

Deleting Maps

If you no longer need a map, you can delete it (presuming that you have edit permission). Deleting a map does not delete any devices or VPNs shown on the map, nor does it delete or modify their configurations; only the map is deleted.

When you delete a map, it is permanently deleted from the server. Other users cannot use the deleted map.

To delete a map, select **Map > Delete Map**, select the map you want to delete from the available maps list and click **OK**. You are asked to confirm the deletion.

You must already be in Map view (select **View > Map View**) to delete a map.

Exporting Maps

When viewing a map, you can export the map to a scalable vector graphics (SVG) image file for use outside of Security Manager.

Related Topics

- Working With Maps, page 35-8
- Understanding Map Elements, page 35-14

Step 1 Select **Map > Export Map**. The Export Topology Map to SVG dialog box opens.

Step 2 Browse to the location in which to save the file.

Step 3 Enter a filename in the File name field. The correct file extension will be added for you.

Step 4 Click **Save**.

Arranging Map Elements

To move a map element, click and hold, then drag it to the desired position. Attached links move automatically, but the other end of the link remains where it is.

You can also automatically arrange the network nodes on the map in several predefined layouts. Only nodes that are already displayed on the map are arranged. Any nodes that you later add do not follow the layout.

To select a map layout, right-click the map background, then select one of the following layouts from the map context menu:

- Hierarchical Layout—Arranges the nodes in a hierarchical layout.
- Radial Layout—Arranges the nodes in a radial layout.
- Circular Layout—Arranges the nodes in a circular layout.

Panning, Centering, and Zooming Maps

There are many options for navigating maps. You can pan the map (move around in the map without changing the zoom level), pan a map so that a particular map element is centered in your view, or zoom in or out to see a different map extent.

To pan a map without changing the zoom level:

- Click the Pan Map toolbar button, then click and hold anywhere on the map and drag the cursor.
- Use the vertical and horizontal scroll bars that are available if the entire map does not fit in the visible page.
- Click and drag the shaded rectangle in the navigation window.
- To center the display of the map on a particular map element, right-click the element, then select **Move to Center**.

To zoom in or out of a map:

- To change the zoom level of the map in predefined increments:
  - To zoom in on the map, select **Map > Zoom In**, or click the **Zoom In** toolbar button.
  - To zoom out from the map, select **Map > Zoom Out**, or click the **Zoom Out** toolbar button.
- To zoom into a specific area of the map, click **Zoom Rectangle** in the map toolbar, then click the map and drag a rectangle around the area. When you release the mouse button, the map zooms to display the area defined by the rectangle.
- Alternatively, to zoom in to or out of a specific area of the map, click and drag the corner of the shaded rectangle in the navigation window.
- To display the entire map, select **Map > Fit to Window**.
- To display the map at actual size, select **Map > Display Actual Size**.

Related Topics

- Using the Navigation Window, page 35-4

Selecting Map Elements

The following table describes how to select map elements. If the selected element contains other elements (for example, a Catalyst switch that contains an FWSM), the containment relationship is shown. For more information, see Showing Containment of Catalyst Switches, Firewalls, and Adaptive Security Appliances, page 35-16.

Table 35-8 Selecting Network Elements

To select...	Do the following
A single map element	Click the element.
Multiple noncontiguous map elements	Ctrl+click each element.
Multiple contiguous map elements	Click the map and drag a rectangle that includes the elements.

Searching for Map Nodes

To search for a map node to help you find it in the active map, select **Map > Find Map Node**. This command opens the Find Node dialog box.

The Find Node dialog box initially lists all objects on the map. Use the fields above the list to filter it (the list shows only objects that satisfy all filter criteria). When you find the desired node, select it in the list and click **OK** to have the node centered and selected in the map.

To filter the list, you can:

- Select a node type from the **Type** list to show only objects of that type.
- Enter the name, or at least the initial characters of the name, in the **Name** field. The list is filtered as you type. Your search term must be from the start of the object name. You cannot use wildcard characters.
- Enter all or part of the IP address or subnet mask. The list is filtered as you enter information.

Using Linked Maps

A linked map is a map that you associate with a map element on another map. Because it is not practical to include all the nodes on a large network in a single map, you can use linked maps to create a hierarchical topology of your network.

You cannot link a node to the another node in the same map.

Before You Begin

You must create the map to link to before you can link to it.

Step 1 Right-click the map element to which to link a map, then select **Set Linked Map**. The Set Linked Map dialog box opens.

Step 2 Select a map to associate with the selected map element, then click **OK**.

Step 3 To open the linked map, right-click the linked node, then select **Open Linked Map**. The current map closes and the linked map opens.

Setting the Map Background Properties

You can change the background of a map by changing the color or by configuring an image. A suggested use for a background image is to use an image that represents a geographic area. Then you can position map elements according to their geographic locations.

Some background images are included with Security Manager. You can also transfer images to the server to use as background images. You can use background images of the following file formats: JPEG, GIF, PNG, IVL, and SVG. If you want to use a new image, copy the image file to the Security Manager server file system by connecting directly to the server. For security reasons, Security Manager does not provide a method of transferring files to the server.

To configure the map background, in Map view, select **Map > Map Properties** to open the Map Settings dialog box.

- To configure a background image, select it in the file list. (Select **none** to remove the map's background image.)

  If the image is not listed, click **Add** and browse to the file you placed on the server using the Import Background Image dialog box. Click **OK** to have Security Manager add it to the list of available background images.

  If you no longer need a listed image, select it and click **Delete**.

Tip You can control the position and scale of the image using the X and Y coordinates and scale settings. The X,Y source point is the upper left corner of the image. You can use positive or negative numbers. You must experiment to get the results you desire. The scale setting is in percentage.

- To change the background color, click **Select** next to the background color field and choose the desired color.

Displaying Your Network on the Map

You use the map view to represent your network topology by creating maps. A map is a visual representation of your network, or a portion of it if it is too large to fit on a single map. Maps consist of map elements that represent devices, links, and other objects in your network. For more information about map, see Working With Maps, page 35-8.

The following topics describe how to create maps:

- Understanding Map Elements, page 35-14
- Displaying Managed Devices on the Map, page 35-16
- Showing Containment of Catalyst Switches, Firewalls, and Adaptive Security Appliances, page 35-16
- Using Map Objects To Represent Network Topology, page 35-17
- Creating and Managing Layer 3 Links on the Map, page 35-19

Understanding Map Elements

All objects that can appear on a map are map elements. You display map elements on a map to create a representation of a portion of your network. For more information about maps, see Working With Maps, page 35-8. To open a map, see Opening Maps, page 35-10.

The following tables describe the elements that can appear on a map:

- Table 35-9 on page 35-14 describes the device nodes that can appear on a map. These elements are managed by Security Manager.
- Table 35-10 on page 35-15 describes the map objects that can appear on a map. These elements are not managed by Security Manager.
- Table 35-11 on page 35-15 describes the map element indicators that can appear with a device node.

Table 35-9 Device Node Types

Node Type	Icon	Description
Firewall		When you select a device, its security contexts are highlighted.
Firewall security context		When you select a security context, the parent device is highlighted. The dotted outline distinguishes the icon as a security context.
Adaptive Security Appliance		When you select a device, its security contexts are highlighted.
Adaptive Security Appliance security context		When you select a security context, the parent device is highlighted. The dotted outline distinguishes the icon as a security context.
Router		Router or VPN concentrator.

Table 35-9 Device Node Types (continued)

Node Type	Icon	Description
Catalyst 6500/7600 or Catalyst switch		When you select a Catalyst device node, any Firewall Service Modules contained in it are highlighted.
Firewall Services Module (FWSM)		When you select a Firewall Services Module, the security contexts it contains are highlighted on the map.
FWSM security context		When you select a security context, the parent device is highlighted. The dotted outline distinguishes the icon as a security context.
IPS Sensor or Security Service Module		An IPS sensor.
VPN connection		Any type of VPN connection. For GET VPNs, a dashed line indicates the connection between group members and key servers.

Table 35-10 Map Object Types

Node Type	Icon	Description
Unmanaged firewall		Unmanaged firewall device.
Unmanaged router		Unmanaged router.
Network		Network with a specified address space.
Host		Network host. Examples: CSA, Syslog Server, CA Server, AAA Host
Cloud		An unspecified group of map objects that provides connectivity between specified nodes.
Layer 3 link	—	Layer 3 network connection

Table 35-11 Map Element Indicators

Indicator	Icon	Description
Linked map		Node is linked to another map.

Related Topics

- Using Map Objects To Represent Network Topology, page 35-17

- Creating and Managing Layer 3 Links on the Map, page 35-19

Displaying Managed Devices on the Map

A device node represents a device that is managed by Security Manager. You add a device node to a map by selecting the device from the Security Manager inventory.

When you add a device node to a map, its layer 3 connectivity to other nodes on the map is created automatically. For more information, see Creating and Managing Layer 3 Links on the Map, page 35-19.

You can add, remove, or show managed nodes by the following means:

- **To add devices that are already in the Security Manager inventory**—Select **Map > Show Devices on Map** to open a device selector. Select the desired devices from the list of available devices and click **>>** to move them to the selected devices list. You can select device groups to move all devices in the group. Click **OK** when the list of selected devices has the desired nodes. Only those devices in the selected list are shown on the map.

  You can remove devices by selecting them in the selected list and clicking **<<**.
- **To add a new device to the map and the device inventory**—Click the **New Device** button in the map toolbar or right-click the map background and select **New Device**. The New Device dialog box opens. Follow the procedures for adding new devices described in Adding Devices to the Device Inventory, page 3-7.
- **To remove a managed node**—Right-click the node and select **Remove from Map**.
- **To locate a device on the open map when in Device view**—Right-click the device in the device selector and select **Show in Map view**. If the device is shown on the active map, it is shown centered and highlighted on the undocked map. You are told that the device cannot be found if the device is not shown on the active map.
- **To locate a device in Device view from the map**—Right-click the device and select **Show in Device View**. Device view is opened with the device selected so that you can edit its policies.

Related Topics

- Understanding Map Elements, page 35-14
- Showing Containment of Catalyst Switches, Firewalls, and Adaptive Security Appliances, page 35-16

Showing Containment of Catalyst Switches, Firewalls, and Adaptive Security Appliances

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any enhancements.

The containment relationship between Catalyst and Adaptive Security Appliance (ASA) devices and their service modules and security contexts, between PIX 7.x+ devices and FWSM and their security contexts, or between IPS devices and their virtual sensors, is displayed in maps as follows:

- When you select a Catalyst device, nodes that represent its Firewall Services Modules (FWSM) are highlighted.

- When you select an ASA, nodes that represent its Security Service Modules are highlighted.
- When you select a service module, the device that contains it is highlighted.
- When you select an IPS device, the nodes that represent virtual sensors defined on the device are highlighted.
- You can view a list of the security contexts contained in an ASA, firewall, or FWSM device, or the virtual sensors contained in an IPS device, by right-clicking the node and selecting **Show Containment**. This command also shows the service modules in a device that has them.
- When you select a security context node, all its ancestor device nodes are highlighted.
- When you select a virtual sensor, the device on which it is defined is highlighted.

Using Map Objects To Represent Network Topology

You can add map elements to a map that represent objects (such as devices and links) that Security Manager does not manage. These nodes are called map objects. You can use map objects to create a more useful representation of your network topology. (If you want to add a managed device, see Displaying Managed Devices on the Map, page 35-16.)

You can add layer 3 links between any map elements, whether they are device nodes, map nodes, or a combination of both types.

Tip To delete a map object, right-click the object and select **Delete Map Object**.

Step 1 Select **Map > Add Map Object**. The Add Map Object dialog box appears (see Add Map Object and Node Properties Dialog Boxes, page 35-18).

Step 2 Do one of the following:

- If you are adding a map object based on the definition of an Security Manager policy object, click **Copy Policy Object** to open the Select Policy Object Dialog Box, page 35-18. Then, select the type of object (AAA server, network/host, PKI enrollment), click **Select** to choose the object, then click **OK** in the Select Policy Object dialog box. Information from the policy object is entered in the Add Map Object dialog box.

  The name of the object is used as the map object name, but you can edit this if desired.
- If you are adding a map object that is not based on a policy object, enter a name for the map object in the **Name** field.

Step 3 Select the type of object that the node represents from the **Type** list. If you selected a policy object, the type is pre-selected, but you can change the selection.

Step 4 (Optional) Add interfaces to the node by doing the following for each interface:

a. Click **Add** to open the Interface Properties Dialog Box, page 35-19. If items already appear in the list, you can select them and click **Edit** to change them.

b. Enter an interface name, IP address, and network mask, then click **OK**.

Step 5 Click **OK**. The map object is added to the center of the map. Move it to the desired location.

Add Map Object and Node Properties Dialog Boxes

For unmanaged map objects, the Add Map Object and Node Properties dialog boxes are the same. Use the Add Map Object dialog box to add an object to the map. Use the Node Properties dialog box to view or edit map object properties. For more information, see Using Map Objects To Represent Network Topology, page 35-17.

For managed map objects (such as a managed device), the Node Properties dialog box is read-only. It displays the object name, type, and list of interface names and IP addresses (if any are defined in Security Manager for the device). The reference information below does not apply to this version of the Node Properties dialog box.

Navigation Path

- To open the Add Map Object dialog box, select **Map > Add Map Object**.
- To open the Node Properties dialog box, right-click a map object or managed device and select **Node Properties**.

Field Reference

Table 35-12 Add Map Object and Node Properties Dialog Boxes for Unmanaged Nodes

Element	Description
Name	The name of the map object. If you select a policy object, the name of the object is automatically used, but you can change it.
Copy Policy Object button	Click to browse for a policy object to use as the basis for the map object using the Select Policy Object Dialog Box, page 35-18.
Type list	The type of object you are creating. If you select a policy object, a type is selected for you, but you can change it if necessary.
Interfaces table	The interfaces on the node. If you select a policy object, information might have been added to this table. • To add an interface, click the **Add (+)** button and fill in the Interface Properties Dialog Box, page 35-19. • To edit an interface, select it and click the **Edit (pencil)** button. • To delete an interface, select it and click the **Delete (trash can)** button.

Select Policy Object Dialog Box

Use the Select Policy Object dialog box to add an object to the map that is defined in a policy object.

Select the type of object that defines the node you want to add to the map from the **Select a Policy Object** list, then click **Select** to select the specific policy object. If you know the object's name, you can type it into the text box instead of clicking Select.

For more information, see Using Map Objects To Represent Network Topology, page 35-17.

Navigation Path

To open this dialog box, click **Copy Policy Object** in the Add Map Object dialog box (see Add Map Object and Node Properties Dialog Boxes, page 35-18).

Interface Properties Dialog Box

Use the Interface Properties dialog box to add and edit interfaces on map objects. For more information, see Using Map Objects To Represent Network Topology, page 35-17.

Navigation Path

To open this dialog box, click the **Add** or **Edit** button in the Add Map Object and Node Properties Dialog Boxes, page 35-18.

Field Reference

Table 35-13 Interface Properties Dialog Box

Element	Description
Interface Name	The interface name.
Interface IP Addr/Mask	The interface IP address and network mask, for example, 10.100.10.0/24 or 10.100.10.0/255.255.255.0.

Creating and Managing Layer 3 Links on the Map

A layer 3 link is a line on the map that represents a network connection between two device interfaces.

Layer 3 connectivity information is automatically added to the map when you add map elements that have interface information. When you add a map element that has interface information, one of the following happens:

- If the interface is on a network that is not represented on the map as a network map object, a network map object is added to the map with a layer 3 link to the new map element.
- If the interface is on a network that is represented on the map as a network map object, a layer 3 link is added between the new map element and the network map object.

When you remove a node interface that is a layer 3 link endpoint, the link is also removed.

You can add additional layer 3 links between device nodes and map objects to illustrate your network's connectivity. Adding Layer 3 links to a map does not configure any network devices. Layer 3 links are just visual elements on the map.

You create layer 3 links to connect any two interfaces on a map. Depending on the interfaces that you choose, the layer 3 link might include intermediary networks or network clouds. In some cases, you have the option to select which intermediary networks and networks clouds are inserted between the connected interfaces.

The following procedure explains how to manually create a new layer 3 link.

Tips

- The automatic addition of network objects and links is called Autolink. You can configure Autolink to not automatically add private or certain reserved network addresses. To configure these settings, select **Tools > Security Manager Administration**, then click **Autolink**.
- To view the properties of a link, right-click the layer 3 link and select **Link Properties**.
- To delete a layer 3 link, right-click the layer 3 link to be removed and select **Delete Link**. Deleting a layer 3 link does not delete any intermediary network or network clouds between map elements.

Step 1 In Map view, click **Map > Add Link** or the Add Link button in the toolbar.

Step 2 Click one of the map elements to connect, then click the other map element to connect.

Step 3 If the map elements contain interfaces, select the source and destination interfaces for the link in the Select Interfaces and Link Properties Dialog Boxes, then click **OK**.

The Add Link dialog box might open, depending on which interfaces you select.

Step 4 If the Add Link Dialog Box opens, select which intermediary objects and network clouds to insert, then click **OK**.

Select Interfaces and Link Properties Dialog Boxes

The Select Interfaces and Link Properties dialog boxes are used with layer 3 links on maps. These dialog boxes show information about the source and destination devices for the link (the source being the first device you clicked when making the link).

If you are creating a link, the Select Interfaces dialog box is used. If there are interfaces defined for the device in Security Manager, select the desired source and destination interfaces for the link you are creating from the **Source/Destination Interface** list.

Tip When creating a link, if there are no interfaces defined for either device, the Interface lists are greyed out. If one device has interfaces defined, both fields are active, but empty for the device that does not have interfaces defined for it. You cannot change the interface when viewing link properties.

Navigation Path

For information on how to create layer 3 links or view their properties, see Creating and Managing Layer 3 Links on the Map, page 35-19.

Add Link Dialog Box

Use the Add Link dialog box to select how to represent the layer 3 link that you are adding to the map.

The contents of the Add Link dialog box vary according to which nodes and interfaces you are connecting. Select the check boxes for each intermediary map object (network or cloud) that you want to insert between the connected nodes. If desired, you can change the names of the map objects.

Navigation Path

This dialog box might open when you add a link between nodes, depending on which interfaces you select to connect. For the procedure, see Creating and Managing Layer 3 Links on the Map, page 35-19.

Managing VPNs in Map View

The following topics describe how to manage VPNs in the Map view:

- Displaying Existing VPNs on the Map, page 35-21
- Creating VPN Topologies in Map View, page 35-21
- Editing VPN Policies or Peers From the Map, page 35-22

Displaying Existing VPNs on the Map

To display an existing VPN on the map, select **Map > Show VPNs on Map**. You are prompted with a list of existing VPNs. Select the ones you want from the available VPNs list and click **>>** to move them to the selected list.

Tip You can also remove a VPN using this command. Select the VPNs you want to remove from the selected VPNs list and click **<<**. When you remove a VPN, only the VPN tunnels are removed. The device nodes remain on the map.

When you display a VPN, all of the its member devices are added to the map as device nodes, and all of its tunnels are highlighted. However, devices that you removed from the map previously are not added, even if they are members of a VPN that you display. You can add such devices to the map manually, and their VPN connectivity is displayed.

A VPN tunnel is a line on the map that represents a VPN connection between two devices. VPN tunnels are not added to the map automatically when you add a device node that is a member of a VPN. However, if the VPN was already selected to be shown on the map, adding a device in the VPN to the map will also display the tunnel.

For an explanation of the icons used in the map, see Understanding Map Elements, page 35-14.

Creating VPN Topologies in Map View

You can create VPN connections between VPN-capable managed device nodes that are displayed on the map. You cannot create Extranet VPNs, however.

To create a VPN, do one of the following:

- Click the **New VPN** button in the toolbar and select the type of VPN you want to configure: point-to-point, hub and spoke, or full mesh.
- Select the devices that you want to participate in the VPN (use Ctrl+click to select multiple devices), and either right click and select the command for the desired type of VPN, or click the **New VPN** button and select the VPN type.

  Consider the following tips:

  - Select only 2 devices to create a point-to-point VPN.
  - If you create a hub-and-spoke VPN, the device you right-click is initially defined as the hub, but you can change that in the wizard.
  - While in the wizard, you can add or remove devices. You are not restricted to the devices you selected on the map.

Using either technique, the Create VPN wizard opens, where you can create the VPN. For more information, see Creating or Editing VPN Topologies, page 25-28 or click the Help button in the wizard.

The VPN is displayed on the map when you are finished with the wizard.

Related Topics

- Selecting Map Elements, page 35-12

Editing VPN Policies or Peers From the Map

You can edit VPN policies, or the peers that participate in a VPN, from map view. To edit policies or peers, right-click a VPN tunnel or device node and select one of these commands:

- **Edit VPN Policies**—To open the Site-to-Site VPN Manager, where you can edit the policies that define the VPN. For more information, see Site-to-Site VPN Manager Window, page 25-18.
- **Edit VPN Peers**—To open a dialog box that allows you to configure the peers that participate in the VPN. Click the **Help** button in the dialog box for more information.
- **Show VPN Peers**—To see which devices participate in a VPN without editing the list (VPN Peers dialog box).

If the device participates in more than one VPN, you are first prompted to select the desired VPN (with the Select VPN to Configure dialog box) before the appropriate dialog box is opened.

Managing Device Policies in Map View

You can perform only basic policy management and configure firewall services policies in Map view. You cannot configure other types of policies. The following topics describe how to manage policies from the Map view:

- Performing Basic Policy Management in Map View, page 35-22
- Managing Firewall Policies in Map View, page 35-23
- Managing Firewall Settings in Map View, page 35-23

Performing Basic Policy Management in Map View

You can perform some basic policy management tasks in Map view. Right click the device and select one of the following commands:

- **Copy Policies Between Devices**—To copy local device policies from one device to another. For more information on copying policies, see Copying Policies Between Devices, page 5-33.
- **Share Device Policies**—To create shared policies from local device policies. For more information on sharing policies, see Sharing Multiple Policies of a Selected Device, page 5-42.
- **Clone Device**—To create a copy of a device, including its policies. For more information on cloning devices, see Cloning a Device, page 3-56.
- **Preview Configuration**—To view the configuration file that will be generated for the device, including the changes from the previous deployment. For more information on previewing configurations, see Previewing Configurations, page 8-44.
- **Discover Policies on Device**—To discover the policies defined on the device and configure them in Security Manager, wiping out whatever policies are defined in Security Manager for the device. For more information device discovery, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Related Topics

- Chapter 8, “Managing Deployment”
- Chapter 3, “Managing the Device Inventory”

- Chapter 5, "Managing Policies"

Managing Firewall Policies in Map View

You can configure firewall policies on a device in Map view. These policies are local to the device rather than being shared policies (you must use Policy view to configure shared policies).

Tip If you want to assign a shared policy to a device, see Performing Basic Policy Management in Map View, page 35-22.

To configure local firewall policies on a device in Map view, right click the device and select one of the following commands:

- **Edit Firewall Policies > AAA Rules**—To configure AAA policies, which control who is allowed access to the device and what services they are allowed to use once they have access. For more information on configuring AAA rules, see AAA Rules Page, page 15-10.
- **Edit Firewall Policies > Access Rules**—To configure Access Rules policies, which control the traffic that flows through a device. For more information on configuring access rules, see Access Rules Page, page 16-10.
- **Edit Firewall Policies > Inspection Rules**—To configure Inspection Rules policies, which analyze traffic at the application layer and track TCP and UDP sessions to perform refined access control. For more information on configuring inspection rules, see Inspection Rules Page, page 17-8.
- **Edit Firewall Policies > Botnet Traffic Filter Rules**—(ASA 8.2 and higher only) To configure Botnet Traffic Filter Rules policies, which monitor web traffic. For more information on configuring botnet traffic filter rules, see Botnet Traffic Filter Rules Page, page 19-9.
- **Edit Firewall Policies > Transparent Rules**—To configure Transparent Rules policies, which define EtherType rules for transparent firewalls. For more information on configuring inspection rules, see Transparent Rules Page, page 23-3.
- **Edit Firewall Policies > Web Filter Rules**—To configure Web Filter Rules policies, which define rules for web access. For more information on configuring web filter rules, see Web Filter Rules Page (ASA/PIX/FWSM), page 18-3 or Web Filter Rules Page (IOS), page 18-12.
- **Edit Firewall Policies > Zone Based Firewall Rules**—(IOS 12.4(6)T and higher only) To configure Zone Based Firewall Rules policies, which configure inspection and web filtering using security zones. For more information on configuring zone based firewall rules, see Zone-based Firewall Rules Page, page 21-58.

Related Topics

- Chapter 12, "Introduction to Firewall Services"
- Chapter 5, "Managing Policies"

Managing Firewall Settings in Map View

You can configure firewall settings policies on a device in Map view. These policies are local to the device rather than being shared policies (you must use Policy view to configure shared policies).

Tip If you want to assign a shared policy to a device, see Performing Basic Policy Management in Map View, page 35-22.

To configure local firewall settings policies on a device in Map view, right click the device and select one of the following commands:

- **Edit Firewall Settings > AAA Firewall**—(ASA/PIX/FWSM only) To configure AAA Firewall settings policies, which configures proxy, authentication challenge, MAC exempt lists, and other general AAA settings. For more information on configuring AAA firewall settings, see AAA Firewall Settings Page, Advanced Setting Tab, page 15-20 and AAA Firewall Page, MAC-Exempt List Tab, page 15-26.
- **Edit Firewall Settings > Access Control**—To configure Access Control settings policies, which configures optimization and other general access control settings. For more information on configuring access control settings, see Access Control Settings Page, page 16-24.
- **Edit Firewall Settings > AuthProxy**—(IOS devices only) To configure AuthProxy settings policies, which configure general settings for authorization proxies. For more information on configuring authorization proxies, see AAA Page, page 15-28.
- **Edit Firewall Settings > Inspection**—(IOS devices only) To configure Inspection settings policies, which configure timeout and session settings for inspection rules. For more information on configuring inspection settings, see Configuring Settings for Inspection Rules for IOS Devices, page 17-111.
- **Edit Firewall Settings > Web Filter**—To configure Web Filter settings policies, which configure the server used for web filtering. For more information on configuring web filter settings, see Web Filter Settings Page, page 18-16.
- **Edit Firewall Settings > Zone Based Firewall**—(IOS 12.4(6)T and higher devices) To configure Zone Based Firewall settings policies, which configure zones and Trend web filter server settings.

PART 4

IPS Configuration

CHAPTER 36

Getting Started with IPS Configuration

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

Cisco Intrusion Prevention System (IPS) Sensors are network devices that perform real-time monitoring of network traffic for suspicious activities and active network attacks. The IPS sensor analyzes network packets and flows to determine whether their contents appear to indicate an attack against your network.

Using Cisco Security Manager, you can configure and manage sensors, which can be dedicated stand-alone network appliances, Catalyst 6500 switch modules, service modules running in supported ASA devices or routers, and IPS-enabled Cisco IOS Software images running on integrated services routers. For a full list of supported IPS devices and software versions, see the Supported Devices and Software Versions for Cisco Security Manager document for this version of the product.

This chapter contains the following topics:

- Understanding IPS Network Sensing, page 36-2
- Overview of IPS Configuration, page 36-5
- Identifying Allowed Hosts, page 36-7
- Configuring SNMP, page 36-8
- Managing User Accounts and Password Requirements, page 36-15
- Identifying an NTP Server, page 36-23
- Identifying DNS Servers, page 36-24
- Identifying an HTTP Proxy Server, page 36-24
- IPS SSHv2 Known Host Keys, page 36-25
- Configuring IPS SSHv1 Fallback Settings, page 36-26
- Configuring the External Product Interface, page 36-26
- Configuring IPS Logging Policies, page 36-30
- IPS Health Monitor, page 36-31
- Configuring IPS Security Settings, page 36-32

Understanding IPS Network Sensing

Network sensing can be accomplished using Cisco IPS sensors (appliances, switch modules, network modules, and SSMs) and Cisco IOS IPS devices (Cisco IOS routers with IPS-enabled images and Cisco ISRs). These sensing platforms are components of the Cisco Intrusion Prevention System and can be managed and configured through Cisco Security Manager. These sensing platforms monitor and analyze network traffic in real time. They do this by looking for anomalies and misuse on the basis of network flow validation, an extensive embedded signature library, and anomaly detection engines. However, these platforms differ in how they can respond to perceived intrusions.

Tip Cisco IPS sensors and Cisco IOS IPS devices are often referred to collectively as IPS devices or simply sensors. However, Cisco IOS IPS does not run the full dedicated IPS software, and its configuration does not include IPS device-specific policies. Additionally, the amount of sensing that you can perform with Cisco IOS IPS is more limited. The following sections focus on using dedicated IPS devices, including service modules installed in IOS routers, rather than Cisco IOS IPS. For a discussion focused on Cisco IOS IPS, see Intrusion Prevention System (IPS) Cisco IOS Intrusion Prevention System Deployment Guide on Cisco.com and Chapter 45, "Configuring IOS IPS Routers". Also, see http://www.cisco.com/go/iosips.

When an IPS device detects unauthorized network activity, it can terminate the connection, permanently block the associated host, and take other actions.

Note For more overview information on IPS sensors, including a comparison of the available appliances and service modules and details about device interfaces, see Introducing the Sensor in *Installing Cisco Intrusion Prevention System Appliances and Modules*. A list of these documents for each IPS release is available at http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/prod_installation_guides_list.html.

This section contains the following topics:

- Capturing Network Traffic, page 36-2
- Correctly Deploying the Sensor, page 36-4
- Tuning the IPS, page 36-4

Capturing Network Traffic

The sensor can operate in either promiscuous or inline mode. The following illustration shows how you can deploy a combination of sensors operating in both inline (IPS) and promiscuous (IDS) modes to protect your network.

Figure 36-1 Comprehensive IPS Deployment Solutions

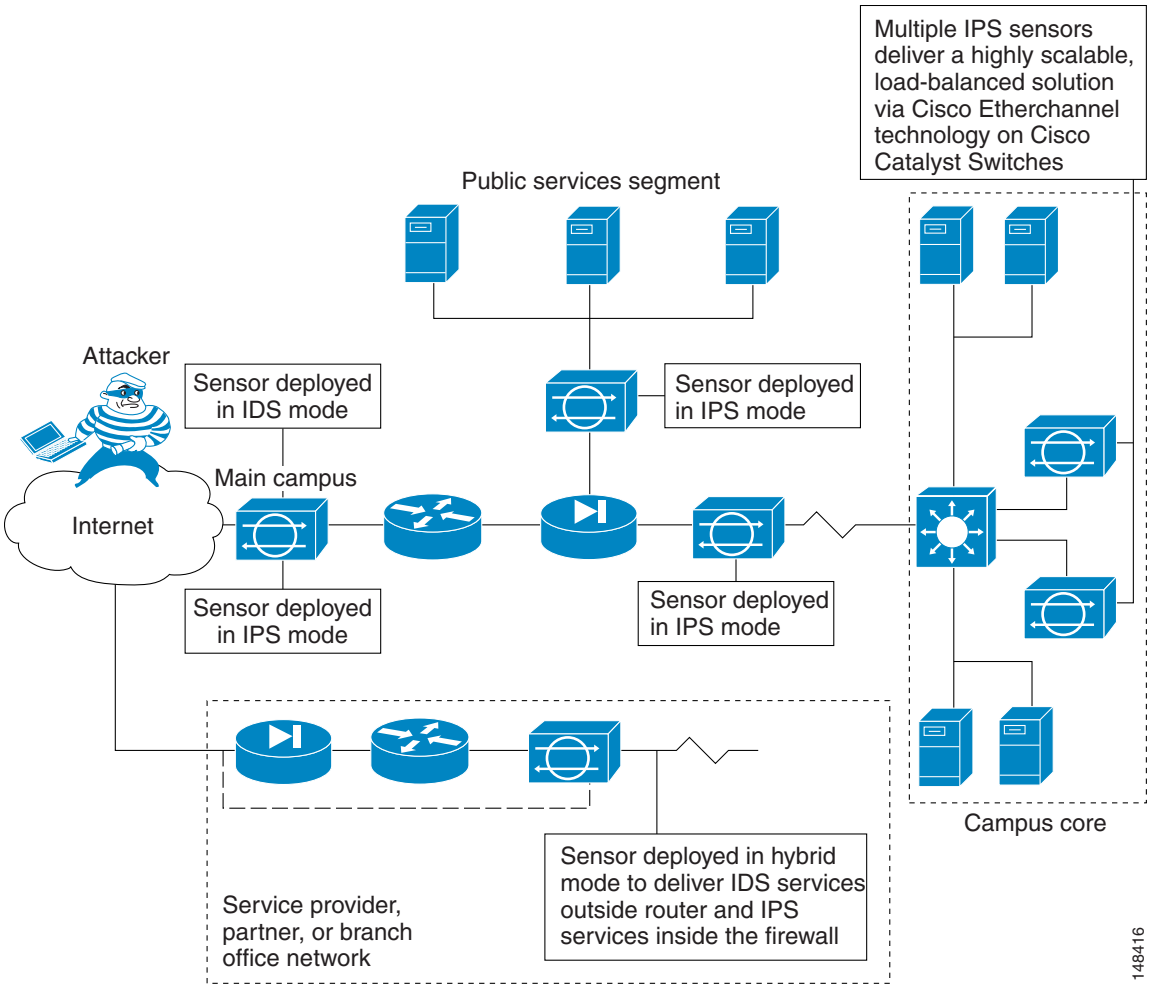

The command and control interface is always Ethernet. This interface has an assigned IP address, which allows it to communicate with the manager workstation or network devices (Cisco switches, routers, and firewalls). Because this interface is visible on the network, you should use encryption to maintain data privacy. SSH is used to protect the CLI and TLS/SSL is used to protect the manager workstation. SSH and TLS/SSL are enabled by default on the manager workstations.

When responding to attacks, the sensor can do the following:

- Insert TCP resets via the sensing interface.

Note You should select the TCP reset action only on signatures associated with a TCP-based service. If selected as an action on non-TCP-based services, no action is taken. Additionally, TCP resets are not guaranteed to tear down an offending session because of limitations in the TCP protocol.

- Make ACL changes on switches, routers, and firewalls that the sensor manages.

Note ACLs may block only future traffic, not current traffic.

- Generate IP session logs, session replay, and trigger packets display.

  IP session logs are used to gather information about unauthorized use. IP log files are written when events occur that you have configured the appliance to look for.

- Implement multiple packet drop actions to stop worms and viruses.

Correctly Deploying the Sensor

Before you deploy and configure your sensors, you should understand the following about your network:

- The size and complexity of your network.
- Connections between your network and other networks, including the Internet.
- The amount and type of traffic on your network.

This knowledge will help you determine how many sensors are required, the hardware configuration for each sensor (for example, the size and type of network interface cards), and how many managers are needed.

You should always position the IPS sensor behind a perimeter-filtering device, such as a firewall or adaptive security appliance. The perimeter device filters traffic to match your security policy thus allowing acceptable traffic in to your network. Correct placement significantly reduces the number of alerts, which increases the amount of actionable data you can use to investigate security violations. If you position the IPS sensor on the edge of your network in front of a firewall, your sensor will produce alerts on every single scan and attempted attack even if they have no significance to your network implementation. You will receive hundreds, thousands, or even millions of alerts (in a large enterprise environment) that are not really critical or actionable in your environment. Analyzing this type of data is time consuming and costly.

Tuning the IPS

Tuning the IPS ensures that the alerts you see reflect true actionable information. Without tuning the IPS, it is difficult to do security research or forensics on your network because you will have thousands of benign events, also known as false positives. False positives are a by-product of all IPS devices, but they occur much less frequently in Cisco IPS devices because Cisco IPS devices are stateful, normalized, and use vulnerability signatures for attack evaluation. Cisco IPS devices also provide risk rating, which identifies high risk events, and policy-based management, which lets you deploy rules to enforce IPS signature actions based on risk rating.

Follow these tips when tuning your IPS sensors:

- Place your sensor on your network behind a perimeter-filtering device.

  Proper sensor placement can reduce the number of alerts you need to examine by several thousands a day.

- Deploy the sensor with the default signatures in place.

  The default signature set provides you with a very high security protection posture. The Cisco signature team has spent many hours on testing the defaults to give your sensor the highest protection. If you think that you have lost these defaults, you can restore them.

- Make sure that the event action override is set to drop packets with a risk rating greater than 90.

  This is the default and ensures that high risk alerts are stopped immediately.

- Filter out known false positives caused by specialized software, such as vulnerability scanner and load balancers by one of the following methods:
  - You can configure the sensor to ignore the alerts from the IP addresses of the scanner and load balancer.
  - You can configure the sensor to allow these alerts and then use Event Viewer to filter out the false positives.
- Filter the Informational alerts.

  These low priority events notifications could indicate that another device is doing reconnaissance on a device protected by the IPS. Research the source IP addresses from these Informational alerts to determine what the source is.
- Analyze the remaining actionable alerts:
  - Research the alert.
  - Fix the attack source.
  - Fix the destination host.
  - Modify the IPS policy to provide more information.

Overview of IPS Configuration

There are a wide variety of devices on which you can configure the Intrusion Prevention System. From a configuration point-of-view, you can separate the devices into two groups: dedicated appliances and service modules (for routers, switches, and ASA devices) that run the full IPS software; and IPS-enabled routers running Cisco IOS Software 12.4(11)T and later (Cisco IOS IPS).

The following procedure is an overview of IPS configuration on dedicated appliances and service modules. For Cisco IOS IPS devices (which does not include IPS service modules installed in a router), see Overview of Cisco IOS IPS Configuration, page 45-4.

Step 1 Install and connect the device to your network. Install the device software and perform basic device configuration. Install the licenses required for all of the services running on the device. The amount of initial configuration that you perform influences what you will need to configure in Security Manager.

Follow the instructions in the *Installing Cisco Intrusion Prevention System Appliances and Modules* document for the IPS version you are using.

Step 2 Add the device to the Security Manager device inventory (see Adding Devices to the Device Inventory, page 3-7).

Tip You can discover router and Catalyst switch modules when adding the device in which the module is installed. For ASA devices, you must add the service module separately.

Step 3 Configure the interfaces as described in Configuring Interfaces, page 37-6. You must enable the interfaces connected to your network for the device to function.

For certain types of service module, there are additional policies to configure:

- Router-hosted service modules—Configure the **IPS Module** interface settings policy on the router. For more information, see IPS Module Interface Settings on Cisco IOS Routers, page 62-22.

- IDSM—Configure the **IDSM Settings** Catalyst platform policy. For more information, see IDSM Settings, page 68-43.
- IPS modules on ASA devices—Configure the **Platform > Service Policy Rules > IPS, QoS, and Connection Rules** policy on the host ASA to specify the traffic that should be inspected. For more information, see About IPS Modules on ASA Devices, page 58-15 and Service Policy Rules Page, page 58-5.

Step 4 Use the **Virtual Sensors** policy to assign interfaces to the virtual sensors, including the base vs0 virtual sensor that exists for all IPS devices. For information about virtual sensor settings and assigning interfaces to a virtual sensor, see Defining A Virtual Sensor, page 38-5.

If the device supports it, and you have a need for it, you can also create user-defined virtual sensors so that a single device acts like multiple sensors. Most of the IPS configuration is done on the parent device, but you can configure unique settings per virtual sensor for signatures, anomaly detection, and event actions. For more information, see Chapter 38, "Configuring Virtual Sensors".

Step 5 Configure basic device access platform policies. These policies determine who can log into the device:

- **AAA**—Configure this policy if you want to use a RADIUS server to control access to the device. You can use AAA control in conjunction with local user accounts defined in the User Accounts policy. See Configuring AAA Access Control for IPS Devices, page 36-21.
- **Allowed Hosts**—The addresses of hosts who are allowed access. Ensure that the Security Manager server is included as an allowed host, or you cannot configure the device using Security Manager. See Identifying Allowed Hosts, page 36-7.
- **SNMP**—Configure this policy if you want to use an SNMP application to manage the device. See Configuring SNMP, page 36-8.
- **Password Requirements**—You can define the acceptable characteristics of a user password. See Configuring User Password Requirements, page 36-20.
- **User Accounts**—The user accounts defined on the device. See Configuring IPS User Accounts, page 36-18.

Step 6 Configure basic server access platform policies. These policies identify the servers to which the device can connect:

- **External Product Interface**—If you use Management Center for Cisco Security Agents, configure this policy to allow the sensor to download host postures from the application. See Configuring the External Product Interface, page 36-26.
- **NTP**—Configure this policy if you want to use a Network Time Protocol server to control the device time. See Identifying an NTP Server, page 36-23.
- **DNS, HTTP Proxy**—The DNS and HTTP Proxy policies are required only if you configure global correlation. They identify a server that can resolve DNS names to IP addresses. Use the HTTP Proxy policy if your network requires the use of a proxy to make Internet connections; otherwise, use the DNS policy. See Identifying DNS Servers, page 36-24 or Identifying an HTTP Proxy Server, page 36-24.

Step 7 Configure the Logging policy if you want non-default logging. See Configuring IPS Logging Policies, page 36-30.

Step 8 Configure IPS signatures and event actions. Event action policies are easier to configure than creating custom signatures, so try to use event action filters and overrides to modify signature behavior before trying to edit specific signatures. For more information, see the following topics:

- Chapter 40, "Configuring Event Action Rules"
- Configuring Signatures, page 39-4

Step 9 If you use any of the Request Block or Request Rate Limit event actions, configure blocking or rate limiting hosts. See Configuring IPS Blocking and Rate Limiting, page 43-7.

Step 10 Configure other desired advanced IPS services. See the following topics:

- Chapter 42, "Configuring Global Correlation"
- Configuring Anomaly Detection, page 41-6

Step 11 Maintain the device:

- Update and redeploy configurations as necessary.
- Apply updated signature and engine packages. For information about checking for updates, applying them, and setting up regular automated updates, see Managing IPS Updates, page 44-4.
- Manage the device licenses. You can update and redeploy licenses, or automate license updates. For more information, see the following topics:
  - Updating IPS License Files, page 44-1
  - Redeploying IPS License Files, page 44-2
  - Automating IPS License File Updates, page 44-3
- Manage the certificates required for SSL (HTTPS) communication. These certificates expire, so you need to regenerate them approximately every 2 years. For information on regenerating certificates and ensuring that the certificates defined on the device are synchronized with those stored in the Security Manager certificate store, see Managing IPS Certificates, page 44-10.

Step 12 Monitor the device:

- Use the Event Viewer application to view alerts generated from the device. You can open Event Viewer from the Launch menu in Configuration Manager or Report Manager, or from the Windows Start menu.
  - For information on using Event Viewer, see Chapter 69, "Viewing Events".
  - For an example of how to filter IPS alerts, see Removing False Positive IPS Events from the Event Table, page 69-66.
- Use the Report Manager application to generate reports on IPS usage, including comparisons of inline vs. promiscuous mode, and global correlation vs. traditional inspection. You can also analyze top attackers, victims, signatures, blocked signatures, and perform target analysis. The following topics explain Report Manager and the IPS reports in more detail:
  - Chapter 70, "Managing Reports"
  - Understanding General IPS Reports, page 70-19
  - Understanding IPS Top Reports, page 70-17
  - Opening and Generating Reports, page 70-20

Identifying Allowed Hosts

Use the Allowed Hosts policy to identify which hosts or networks have permission to access the IPS sensor. By default, no hosts are permitted to access a sensor, so you must add hosts or networks to this policy.

Specifically, you must add either the IP address of the Security Manager server, or its network address, or Security Manager cannot configure the device. Also add the addresses of all other management hosts that you use, such as CS-MARS.

Tip If you add host addresses only, you will be limited to using those workstations to access the device. Instead, you can specify network addresses to allow all hosts connected to specific "safe" networks access.

Step 1 Do one of the following to open the Allowed Hosts policy:

- (Device view) Select **Platform > Device Admin > Device Access > Allowed Hosts** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Allowed Hosts**, then select an existing policy or create a new one.

Step 2 Do one of the following:

- To add an entry, click the **Add Row** button and fill in the Access List dialog box.

  You can add up to 512 entries.
- To edit an entry, select it and click the **Edit Row** button.
- To delete an entry, select it and click the **Delete Row** button.

Step 3 When adding or editing an entry, specify the host or network address in the Add or Modify Access List dialog box, then click **OK**. You can enter addresses using the following formats:

- Host address—A simple IP address, such as 10.100.10.10.
- Network address—A network address and mask, such as 10.100.10.0/24 or 10.100.10.0/255.255.255.0.
- A network/host policy object—Click **Select** to select an existing object or to create a new one. To use the object in this policy, it must have a single value, either a single network or a single host.

Configuring SNMP

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

SNMP is a simple request/response protocol. The network-management system issues a request, and managed devices return responses. This behavior is implemented by using one of four protocol operations: Get, GetNext, Set, and Trap.

You can configure the sensor for monitoring by SNMP. SNMP defines a standard way for network management stations to monitor the health and status of many types of devices, including switches, routers, and sensors.

You can configure the sensor to send SNMP traps. SNMP traps enable an agent to notify the management station of significant events by way of an unsolicited SNMP message.

Trap-directed notification has the following advantage—if a manager is responsible for a large number of devices, and each device has a large number of objects, it is impractical to poll or request information from every object on every device. The solution is for each agent on the managed device to notify the manager without solicitation. It does this by sending a message known as a trap of the event.

After receiving the event, the manager displays it and can take an action based on the event. For example, the manager can poll the agent directly, or poll other associated device agents to get a better understanding of the event.

Tip Trap-directed notification results in substantial savings of network and agent resources by eliminating frivolous SNMP requests. However, it is not possible to totally eliminate SNMP polling. SNMP requests are required for discovery and topology changes. In addition, a managed device agent cannot send a trap if the device has had a catastrophic outage.

This procedure describes how to configure SNMP on an IPS sensor so that you can manage the sensor with an SNMP management station, including the configuration of traps.

Step 1 Do one of the following to open the SNMP policy:

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Device Access > SNMP**, then select an existing policy or create a new one.

Step 2 On the **General Configuration** tab, configure at least the following options. For a complete description of all available options, see General SNMP Configuration Options, page 36-10.

- **Enable SNMP Gets/Sets**—Select this option to enable the SNMP management workstation to obtain (get) information, and to modify (set) values on the IPS sensor. If you do not enable this option, the management workstation cannot manage this sensor.
- **Read-Only Community String**—The community string required for read-only access to the sensor. SNMP get requests from the management station must supply this string to get responses from the sensor. This string gives access to all SNMP get requests.
- **Read-Write Community String**—The community string required for read-write access to the sensor. SNMP set requests from the management station must supply this string to get responses from the sensor; it can also be used on get requests. This string gives access to all SNMP get and set requests.

Step 3 On the **SNMPv3 Users** tab, add one or more SNMPv3 users to configure SNMPv3 settings on the managed IPS devices. Beginning with version 4.6, Security Manager enables you to configure SNMPv3 settings on the IPS devices it manages. For more information, see SNMPv3 Users Tab, page 36-11.

Note SNMPv3 is supported in IPS version 7.2.2 and later, but not in the IPS version 7.3.1. Security Manager can however manage IPS 7.3.1 devices. If you try to use Security Manager to upgrade an IPS device from version 7.2.2, with SNMP policy configured, to version 7.3.1, a mouse-over tooltip displays the message "Selected upgrade is not recommended. Unassign the SNMP policy on the device and deploy it to continue with the upgrade to 7.3.1". For information about managing SNMPv3 policies on IPS devices with version 7.2.2 and later, see the Release Notes for Cisco Intrusion Prevention System 7.2(2).

Step 4 If you want to configure SNMP traps, click the **SNMP Trap Configuration** tab and configure at least the following options. For a complete description of all available options, see SNMP Trap Configuration Tab, page 36-13.

- **Enable Notifications**—Select this option to allow the sensor to send SNMP traps.
- **Trap Destinations**—Add the SNMP management stations that should be trap destinations. Click the **Add Row (+)** button to add a new destination, or select a destination and click the **Edit Row (pencil)** button to change its configuration.

  When adding or editing a trap destination, the trap community string that you enter overrides the default community string entered on the SNMP Trap Configuration tab. The community string appears in the traps sent to this destination and is useful if you are receiving multiple types of traps from multiple agents. For example, a router or sensor could be sending the traps, and if you put something that identifies the router or sensor specifically in your community string, you can filter the traps based on the community string.

  To remove a destination, select it and click the **Delete Row (trash can)** button.

Step 5 If you configure trap destinations, you must also ensure that the desired alerts include the **Request SNMP Trap** action. You have the following options for adding this action:

- (Easy way.) Create an event action override to add the Request SNMP Trap action to all alerts of a specified risk rating (**IPS > Event Actions > Event Action Overrides** policy). For example, you could generate traps for all alerts with a risk rating between 85-100. Event action overrides let you add an action without individually editing each signature. For more information, see Configuring Event Action Overrides, page 40-13.
- (Precise way.) Edit the Signatures policy (**IPS > Signatures > Signatures**) to add the Request SNMP Trap action to the signatures for which you want to send trap notifications. Traps are sent only for signatures that you configure to send traps.

Note If the signature has Default for the source, you have to change the source to the Local source before you can change the action. However, if you right-click the Action cell in the signatures table and select **Edit Actions**, then select Request SNMP Trap (along with any other desired action) and click **OK**, the source is automatically changed to Local.

Step 6 Add the SNMP management stations to the Allowed Hosts policy. The management stations must be allowed hosts to access the sensor. See Identifying Allowed Hosts, page 36-7.

General SNMP Configuration Options

Use the General Configuration tab on the SNMP page to configure general SNMP parameters and apply them to IPS sensors. For the procedure, see Configuring SNMP, page 36-8.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector. Select the General Configuration tab.
- (Policy view) Select **IPS > Platform > Device Admin > Device Access > SNMP**, then select an existing policy or create a new one. Select the General Configuration tab.

Field Reference

Table 36-1 General Configuration Tab, SNMP Policy for IPS Sensors

Element	Description
Enable SNMP Gets/Sets	Whether to enable the SNMP management workstation to obtain (get) information, and modify (set) values on the IPS sensor. If you do not enable this option, the management workstation cannot manage this sensor; the sensor will not respond to SNMP requests.
Read-Only Community String	The community string required for read-only access to the sensor. SNMP get requests from the management station must supply this string to get responses from the sensor. This string gives access to all SNMP get requests. Use the string to help identify the sensor.
Read-Write Community String	The community string required for read-write access to the sensor. SNMP set requests from the management station must supply this string to get responses from the sensor; it can also be used on get requests. This string gives access to all SNMP get and set requests. Use the string to help identify the sensor.
Sensor Contact	The network administrator or contact point who is responsible for this sensor.
Sensor Location	The physical location of the sensor, such as building address, name, and room number.
Sensor Agent Port	The port to use for SNMP get/set communication with the sensor. The default is 161. The valid range is 1 to 65535. Enter a port number or the name of a port list object, or click **Select** to select a port list object from a list or to create a new object. The port list object must identify a single port.
SNMP Agent Protocol	The protocol you are using for SNMP, either UDP (the default) or TCP. Select the protocol used by your SNMP management station.

SNMPv3 Users Tab

Beginning with version 4.6, Security Manager enables you to configure SNMPv3 settings on the IPS devices it manages. You must add SNMPv3 users to configure SNMPv3 settings on the managed IPS devices.

You can use the SNMPv3 Users tab on the SNMP page to view, add, edit, or delete SNMPv3 users.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector. Select the SNMPv3 Users tab.
- (Policy view) Select **IPS > Platform > Device Admin > SNMP**, then select an existing policy or create a new one. Select the SNMPv3 Users tab.

Do one of the following:

- To add an SNMPv3 user, click the **Add Row (+)** button. This opens the Add SNMPv3 User dialog box. Enter the information required to create the user. For detailed information on the settings, see Add SNMPv3 User Dialog Box, page 36-12

- To edit an SNMPv3 user, select it and click the **Edit Row (pencil)** button and make the required changes in the Edit SNMPv3 User dialog box.
- To delete an existing SNMPv3 user, select it and click the **Delete Row (trash can)** button.

Field Reference

Table 36-2 SNMPv3 Users

Element	Description
User Name	Name of the user on the host that belongs to the SNMP agent.
Access Control	Access privilege for the SNMPv3 user.
Security Level	Security level for the SNMPv3 user.
Authentication Protocol	The authentication protocol keyword is the authentication level used to configure the SNMPv3 user.
Privacy Protocol	The privacy protocol keyword is the privacy or encryption algorithm used to configure the SNMPv3 user.

Add SNMPv3 User Dialog Box

Use the Add SNMPv3 User dialog box to configure a new SNMPv3 user for the managed IPS device.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector. Select the SNMPv3 Users tab and click the **Add Row (+)** button.
- (Policy view) Select **IPS > Platform > Device Admin > SNMP**, then select an existing policy or create a new one. Select the SNMPv3 Users tab and click the **Add Row (+)** button.

Field Reference

Table 36-3 Add SNMPv3 Users Dialog Box

Element	Description
User Name	Enter a name for the new SNMPv3 user.
Access Control	Select the access privilege for the new SNMPv3 user.
Security Level	Select one of the following security levels for the SNMPv3 user: • **NoAuthNoPriv**—There is no authentication and no privacy, which means that no security is applied to the messages. • **AuthNoPriv**—There is authentication but there is no privacy, which means that messages are authenticated. • **AuthPriv**—Authentication and privacy are configured, which means that messages are authenticated and encrypted.
Authentication Protocol	Select the authentication protocol keyword that specifies which authentication level should be used. There is no default value.
Privacy Protocol	Select the privacy protocol keyword that specifies which privacy or encryption algorithm should be used. For the encryption algorithm, you can specify the AES keyword. There is no default value.

Table 36-3 Add SNMPv3 Users Dialog Box (continued)

Element	Description
Authentication Passphrase	Enter the authentication passphrase argument that specifies the authentication user password. This password must not be less than eight characters. There is no default value.
Privacy Passphrase	Enter the privacy passphrase argument that specifies the encryption user password. This password must not be less than eight characters. There is no default value.

SNMP Trap Configuration Tab

Use the SNMP Trap Communication tab on the SNMP page to configure traps and apply them to sensors and to identify recipients that the traps should be sent to. For the procedure, see Configuring SNMP, page 36-8.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector. Select the SNMP Trap Configuration tab.
- (Policy view) Select **IPS > Platform > Device Admin > Device Access > SNMP**, then select an existing policy or create a new one. Select the SNMP Trap Configuration tab.

Field Reference

Table 36-4 SNMP Trap Configuration Tab, SNMP Policy for IPS Sensors

Element	Description
Enable Notifications	Whether to enable the sensor to send trap notifications to the trap destinations whenever a specific type of event occurs in a sensor. If you do not select this option, the sensor does not send traps. **Tip** To have the sensor send SNMP traps, you must also select **Request SNMP Trap** as the event action when you configure signatures. Traps are sent only for signatures that you configure to send traps.
Error Filter	The type of events that will generate SNMP traps based on the severity of the event: fatal, error, or warning. Select all severities that you want; use Ctrl+click to select multiple values. The sensor sends notifications of events of the selected severities only.
Enable Detail Traps	Whether to include the full text of the alert in the trap. If you do not select this option, sparse mode is used. Sparse mode includes less than 484 bytes of text for the alert.
Default Trap Community String	The community string used for the traps if no specific string has been set for the trap destination in the Trap Destinations table. **Tip** All traps carry a community string. By default, all traps that have a community string identical to that of the destination are taken by the destination. All other traps are discarded by the destination. However, you can configure the destination to determine which trap strings to accept.

Table 36-4 SNMP Trap Configuration Tab, SNMP Policy for IPS Sensors (continued)

Element	Description
Trap Destinations table	The SNMP management stations that will be sent trap notifications. The table shows the IP address of the management station, the community string added to traps from this sensor, and the port to which traps are sent. • To add a destination, click the **Add Row** button and fill in the Add SNMP Trap Communication dialog box (see SNMP Trap Communication Dialog Box, page 36-14). • To edit a destination, select it, click the **Edit Row** button and make your changes. • To delete a destination, select it and click the **Delete Row** button.

SNMP Trap Communication Dialog Box

Use the Add or Modify SNMP Trap Communication dialog box to configure SNMP trap destinations. These are the SNMP management stations that should receive traps from the IPS sensor.

Navigation Path

Go to the IPS **Platform > Device Admin > Device Access > SNMP** policy, select the **SNMP Trap Configuration** tab, and click the **Add Row** button beneath the Trap Destinations table, or select a destination in the table and click the **Edit Row** button. For more information, see SNMP Trap Configuration Tab, page 36-13.

Field Reference

Table 36-5 SNMP Trap Communication Dialog Box

Element	Description
IP Address	The IP address of the SNMP management station that should receive trap notifications. Enter the IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new object. The network/host object must specify a single host IP address.
Trap Community String	The community string of the trap. If you do not enter a trap string, the default trap string defined on the SNMP Trap Communication tab is used for traps sent to this destination.
Trap Port	The port used by the SNMP management station to receive traps. Enter the port number or the name of a port list object, or click Select to select the object from a list or to create a new one. The port list object must identify a single port.
SNMPv3 User	Enter the username of the SNMPv3 user that you configured by using the Add SNMPv3 User Dialog Box, page 36-12. Leave this field blank if you do not want to associate any SNMPv3 user. **Note** If you enter a username that is not a configured SNMPv3 user, you will receive an error message while trying to save the SNMP trap communication settings. () Also, note that you can add up to a maximum of 23 SNMPv3 users.

Managing User Accounts and Password Requirements

You can configure user accounts and passwords, and general password requirements, for your IPS devices. You can configure local users (defined directly on the device), use a RADIUS AAA server, or use them both in conjunction. The policies used are the **AAA**, **User Accounts**, and **Password Requirements** policies in the **Platform > Device Admin > Device Access** folder.

When you create or edit a local user account in Security Manager, the password you enter must satisfy the requirements defined in the Password Requirements policy. This ensures that new passwords meet your security requirements.

Tip If you change the password requirements, and then make changes to any local user account, the new requirements must be met by all user accounts that have passwords managed by Security Manager. This is because Security Manager reconfigures the passwords for all managed accounts if any single account needs to be reconfigured.

The User Accounts policy allows you to centrally manage the local user accounts for your IPS devices. Using a shared policy can help you ensure that all IPS devices contain the same accounts with the same passwords. However, it is important to understand that passwords are encrypted, so Security Manager cannot discover the actual passwords defined on the device. Security Manager manages the passwords for an account only if you define that password in Security Manager. Security Manager does not manage any user accounts defined in a RADIUS AAA server.

The following topics describe IPS user accounts, and Security Manager discovery and deployment considerations, in more detail:

- Understanding IPS User Roles, page 36-15
- Understanding Managed and Unmanaged IPS Passwords, page 36-16
- Understanding How IPS Passwords are Discovered and Deployed, page 36-17
- Configuring IPS User Accounts, page 36-18
- Configuring User Password Requirements, page 36-20
- Configuring AAA Access Control for IPS Devices, page 36-21

Understanding IPS User Roles

There are four user roles for IPS user accounts:

- **Viewer**—Users can view the device configuration and events, but they cannot modify any configuration data except their user passwords.
- **Operator**—Users can view everything and they can modify the following options:
  - Signature tuning (priority, disable or enable).
  - Virtual sensor definition.
  - Managed routers.
  - Their user passwords.
- **Administrator**—Users can view everything and they can modify all options that Operators can modify in addition to the following:
  - Sensor addressing configuration.

  - List of hosts allowed to connect as configuration or viewing agents.
  - Assignment of physical sensing interfaces.
  - Enable or disable control of physical interfaces.
  - Add and delete users and passwords.
  - Generate new SSH host keys and server certificates.
- **Service**—Only one user with service privileges can exist on a sensor. The service user cannot log in to IDM or IME. The service user logs in to a bash shell rather than the CLI. The service role is a special role that allows you to bypass the CLI if needed.

Note The purpose of the Service account is to provide Cisco Technical Support access to troubleshoot unique and unusual problems. It is not needed for normal system configuration and troubleshooting. You should carefully consider whether you want to create a service account. The service account provides shell access to the system, which makes the system vulnerable. However, you can use the service account to create a password if the administrator password is lost. Analyze your situation to decide if you want a service account existing on the system.

Understanding Managed and Unmanaged IPS Passwords

Every IPS local user account has a password, which allows secure user login to the device. These user passwords are encrypted on the IPS device. Thus, when you add an IPS device to the Security Manager inventory, Security Manager cannot read the actual user passwords.

Because Security Manager cannot read the password, it is unable to deploy newly-discovered user account passwords to the device. To avoid putting user accounts into a state where the passwords are unknown and unusable, Security Manager marks discovered user account passwords as **unmanaged**. The status of a password is indicated in the **Is Password Managed?** column of the **Platform > Device Admin > Device Access > User Accounts** policy:

- If **No** is indicated, the password for this account is not configured in Security Manager. When you deploy this policy, Security Manager will not attempt to configure the password for this user account.
- If **Yes** is indicated, the password for this account was configured or updated in Security Manager. When you deploy this policy, Security Manager reconfigures the passwords for all managed accounts, not just the passwords that changed since the last deployment.

  Because Security Manager configures even unchanged passwords, all managed passwords must satisfy the password requirements defined in the Password Requirements policy.

Thus, you can have a mix of managed and unmanaged account passwords. For example, you can have a set of shared user accounts that are centrally managed, and manage these account passwords in Security Manager. Other accounts might be unique to individuals; if you never edit these account passwords in Security Manager, the user can manage these passwords individually on the device.

Tip If you do not want to manage any user accounts in Security Manager, ensure that the User Accounts policy is empty, or simply unassign the policy (right-click the policy and select **Unassign Policy**). Security Manager will not modify user account configurations.

Understanding How IPS Passwords are Discovered and Deployed

Because user passwords are encrypted on IPS devices, Security Manager has to handle them with special care when discovering policies on the device or deploying configurations. When discovering or deploying user accounts on IPS devices, Security Manager does the following:

- **Discovery**—When you add an IPS device to the inventory, or rediscover policies on it, Security Manager determines the current status of each user account, updates the User Account policy with each discovered username and associated role, and marks the user password as unmanaged (as described in Understanding Managed and Unmanaged IPS Passwords, page 36-16).

  You cannot view the account status through Security Manager, because it is dynamic and can change. However, the Discovery Status window displays the status at discovery. Accounts can have these statuses:

  - **Active**—This state indicates that the account is available for use. Active accounts can be accessed using an authentication token if one has been assigned to the account.
  - **Expired**—This state indicates that the account's authentication token has expired and the account can not be accessed using a token until the token has been updated.
  - **Locked**—This state indicates that logins to the account have been disabled due to too many failed authentication attempts. You should update the password for these accounts.

- **Deployment**—You are warned if any deployed user accounts are in the Expired or Locked state. Any unmanaged passwords are not deployed to the device. Also, keep in mind the following points:

  - If you make changes to any user account on the device, all user accounts with managed passwords are reconfigured. If you also changed the Password Requirements policy, all passwords are compared to the new policy and must meet the new requirements.
  - If you change the password of the user account you defined in the device's properties for Security Manager to use when configuring the device, after successful deployment, Security Manager updates the password in the device properties to the new password. You do not need to manually update the password. To see device properties, select **Tools > Device Properties**.

    This behavior assumes that you selected **Security Manager Device Credentials** for the Connect to Device Using option on the **Tools > Security Manager Administration > Device Communication** page. If you are using the logged-in users credentials for deployment, after successful deployment, the overall deployment is marked as failed, and a message explains how to reestablish connection. See Device Communication Page, page 11-21.
  - If you use out-of-band change detection, changes to passwords are not detected. However, changes to usernames and roles are detected. For more information about out-of-band change detection, see Detecting and Analyzing Out of Band Changes, page 8-45.
  - When previewing configurations, you can see changes to the user accounts by selecting to IPS(Delta – User Passwords). However, passwords are masked. For more information, see Previewing Configurations, page 8-44.
  - If you are rolling back configurations, the user accounts are never rolled back. The current status and configuration of user accounts does not change.

Tip The IPS sensor can accept public keys for RSA authentication when logging into the device through an SSH client. Each user has an associated list of authorized keys. Users can use these keys instead of passwords. Security Manager ignores these keys during discovery and deployment. Thus, if keys are configured, Security Manager does not remove the configuration.

Related Topics

- Discovering Policies, page 5-12
- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34
- Understanding Configuration Rollback, page 8-63
- Understanding Rollback for IPS and IOS IPS, page 8-66

Configuring IPS User Accounts

Use the User Accounts policy to configure local user accounts for IPS devices. Users can use these accounts to log into the device. You can create new users, modify user privileges and passwords, and delete users.

The user accounts policy should have at least these accounts:

- cisco—An account named "cisco" must exist on the device and you cannot delete it.
- An administrator account that Security Manager can use—Security Manager must be able to log into the device to configure it. Typically, you create an account for this purpose. However, you have the option of having Security Manager use the user account of the person deploying configurations to log into the device. You can configure this using the **Connect to Device Using** option on the **Tools > Security Manager Administration > Device Communication** page. See Device Communication Page, page 11-21.

IPS user account configuration is more complicated than it seems. Before you configure IPS user accounts, read the following topics:

- Managing User Accounts and Password Requirements, page 36-15
- Understanding IPS User Roles, page 36-15
- Understanding Managed and Unmanaged IPS Passwords, page 36-16
- Understanding How IPS Passwords are Discovered and Deployed, page 36-17
- Configuring User Password Requirements, page 36-20
- Configuring AAA Access Control for IPS Devices, page 36-21

Tips

- Cisco IOS IPS devices use the same user accounts that are defined for the router. This procedure does not apply to Cisco IOS IPS configurations.
- If you change the password for the user defined in the device properties, which Security Manager uses to deploy configurations to the device, Security Manager uses the existing credentials defined in the device properties to log into the device and deploy changes. After successful deployment, the device properties are then changed to use your new settings. For more information on credentials in device properties, see Device Credentials Page, page 3-45.

Related Topics

- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49

Step 1 Do one of the following to open the User Accounts policy:

- (Device view) Select **Platform > Device Admin > Device Access > User Accounts** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Device Access > User Accounts**, then select an existing policy or create a new one.

The policy shows existing user accounts, including the username, role, and whether the password is managed by Security Manager (as explained in Understanding Managed and Unmanaged IPS Passwords, page 36-16).

Step 2 Do one of the following:

- To add a user account, click the **Add Row (+)** button. This opens the Add User dialog box. Enter the information required to define the account. For detailed information on the settings, see Add User and Edit User Credentials Dialog Boxes, page 36-19.
- To edit a user account, select it and click the **Edit Row (pencil)** button and make the required changes in the Edit User dialog box.

  You cannot change a user role to or from the Service role.
- To delete a user account, select it and click the **Delete Row (trash can)** button. You cannot delete the account named cisco.

Tip All password changes must meet the requirements of the Password Requirements policy. If you change the requirements policy, all new user accounts, or edited accounts, are tested against the new requirements. Although the passwords for existing unedited user accounts are not tested, they too must meet the password requirements if you change any user account defined in this policy, because Security Manager will deploy all of the accounts during the next configuration deployment. Passwords are checked for conformity when you validate policies, which typically happens when you submit changes to the database. For more information, see Understanding How IPS Passwords are Discovered and Deployed, page 36-17.

Add User and Edit User Credentials Dialog Boxes

Use the Add User or Edit User Credentials dialog boxes to add or edit IPS device user accounts.

Navigation Path

From the IPS platform User Accounts policy, click the **Add Row (+)** button to create a new account, or select an existing account and click the **Edit Row (pencil)** button. For information on accessing the User Accounts policy, see Configuring IPS User Accounts, page 36-18.

Field Reference

Table 36-6 Add or Edit User Dialog Box

Element	Description
User Name	The username for the account. The name can be 1 to 64 characters, including uppercase and lowercase letters and numbers, plus the special characters ( ) + : , _ / -] + $. You cannot change the username when editing an account.

Table 36-6 Add or Edit User Dialog Box (continued)

Element	Description
Password Confirm	The password for this user account. Enter the password in both fields. The password must conform to the Password Requirements policy for IPS devices; see Configuring User Password Requirements, page 36-20.
Role	The role for this user. For an explanation of these roles, see Understanding IPS User Roles, page 36-15. **Tip** When editing a user account, you cannot select the Service role. When editing an account assigned to the Service role, you cannot change the role.

Configuring User Password Requirements

Use the IPS platform Password Requirements policy to configure the rules for passwords for local IPS device user accounts. All user-created sensor passwords must conform to the requirements defined in this policy. You can configure password requirements for sensor running IPS software version 6.0 or higher.

Tip The requirements you define here determine what is considered an acceptable password in the User Accounts policy (see Configuring IPS User Accounts, page 36-18). If you change this policy, it can be applied even to unchanged user accounts. For more information about the implications of deploying changes to this policy, see Understanding How IPS Passwords are Discovered and Deployed, page 36-17.

To configure IPS password requirements, select one of the following policies:

- (Device view) Select **Platform > Device Admin > Device Access > Password Requirements** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Password Requirements** from the Policy Type selector, then select an existing policy or create a new one.

The following table explains the password requirement options that you can configure.

Table 36-7 Password Requirements Policy

Element	Description
Attempt Limit	How many times a user is allowed to try to log into the device before you lock the user account due to excessive failed attempts. The default is 0, which indicates unlimited authentication attempts. For security purposes, you should change this number.

Table 36-7 Password Requirements Policy (continued)

Element	Description
Size Range	The minimum and maximum size allowed for user passwords; separate the minimum and maximum with a hyphen. The range is 6 to 64 characters; the default is 8-64. **Tip** If you configure non-zero values for any of the minimum characters options, the minimum size you enter in the Size Range field must be equal to or greater than the sum of those values. For example, you cannot set a minimum password size of eight and also require that passwords must contain at least five lowercase and five uppercase characters.
Minimum Digit Characters	The minimum number of numeric digits that must be in a password.
Minimum Uppercase Characters	The minimum number of uppercase alphabet characters that must be in a password.
Minimum Lowercase Characters	The minimum number of lowercase alphabet characters that must be in a password.
Minimum Other Characters	The minimum number of non-alphanumeric printable characters that must be in a password.
Number of Historical Passwords	The number of historical passwords that you want the sensor to remember for each account. Any attempt to change the password of an account fails if the new password matches any of the remembered passwords. If you specify 0, no previous passwords are remembered.

Configuring AAA Access Control for IPS Devices

Use the AAA policy to configure AAA access control for your IPS devices. The device must use IPS Software release 7.0(4) or above or 7.1.3 or above to configure AAA; for example, neither 7.1.1 nor 7.1.2 supports AAA.

You can configure the IPS device to use a RADIUS AAA server to authenticate user access to the device. By configuring AAA, you can reduce the number of local users defined on the device and take advantage of your existing RADIUS setup. If you configure a AAA server, you can configure the device to allow local user accounts as a fallback mechanism if the RADIUS servers are unavailable.

When configuring AAA, you identify the RADIUS server using a AAA server policy object. You can create the object while configuring the policy, or you can create it in the Policy Object Manager. When you configure the AAA server object, you must adhere to the following restrictions:

- **Host**—You must specify the IP address; you cannot use a DNS name.
- **Timeout**—If you enter a timeout value, it must be from 1 to 512 seconds. The generic AAA server object allows higher numbers, but IPS has a more limited timeout range. The default is 3.
- **Protocol**—RADIUS is the only supported protocol.
- **Key**— You must specify the shared secret key that is defined on the RADIUS server. Although this field is optional for a generic AAA server object, IPS requires a key.
- **Port**—Ensure that the RADIUS Authentication/Authorization port is correct. Note that the default port in the AAA server object is different from the IPS default, which is 1812. You will need to change the port if you want to use the IPS default.

For more information about configuring AAA server objects, see Creating AAA Server Objects, page 6-32.

Tip You must ensure that the user account configured in the device properties exists in the RADIUS server or as a local user account, depending on the authorization method that you use. If you switch between local and AAA modes, or change AAA servers, you must ensure that the account is defined in whatever user account database you are using. If you are using AAA with local fallback, the account should be defined in all databases. This account must exist, with the same password defined in the Security Manager device properties for the device, or deployment to the device will fail. The user account used for discovery and deployment must have administrator privileges.

Related Topics

- Managing User Accounts and Password Requirements, page 36-15
- Configuring IPS User Accounts, page 36-18

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > AAA** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > AAA**, then select an existing policy or create a new one.

Step 2 Configure the following basic properties:

- **Authentication Mode**—Whether to use Local or AAA mode. Local mode uses user accounts defined on the IPS device only. With AAA mode, the RADIUS servers are the primary means of user authentication, and you can configure local user accounts as a fallback mechanism. The default is Local. You must select AAA to configure any other options in this policy.
- **Primary RADIUS Server, Secondary RADIUS Server**—The main (primary) AAA server and a backup server, if any. Enter the name of the AAA server policy object that identifies the RADIUS server, or click **Select** to select it from a list of objects or to create a new object.

  When authenticating users, the IPS device sends the user authentication attempt to the primary server. The secondary server is contacted only if the request to the primary server times out.

Step 3 Configure the following optional properties if you want non-default values:

- **Console Authentication**—How you want to authenticate users who access the IPS device through the console:
  - Local—Users connected through the console port are authenticated through local user accounts.
  - Local and RADIUS—Users connected through the console port are authenticated through RADIUS first. If RADIUS fails, local authentication is attempted.
  - RADIUS—Users connected through the console port are authenticated by RADIUS. If you also select Enable Local Fallback, then users can also be authenticated through the local user accounts.
- **RADIUS NAS ID**—The Network Access ID, which identifies the service requesting authentication. The value can be no NAS-ID, cisco-ips, or a NAS-ID already configured on the RADIUS server. The default is cisco-ips.
- **Enable Local Fallback**—Whether you want to fall back to local user account authentication if all RADIUS servers are unavailable. This option is selected by default. Note that local authentication is not attempted if the RADIUS server responds negatively to the logon attempt; local authentication is tried only if no response is received from the RADIUS server.

- **Default User Role**—The role to assign to users who do not have a role assigned in the RADIUS server. You can make Viewer, Operator, or Administrator the default roles, but not Service; select Unspecified to assign no default role (this is the default). For an explanation of user roles, see Understanding IPS User Roles, page 36-15.

Note User role configuration is very important. If you do not assign a role to the user, either through the default user role or in the RADIUS server, the sensor prevents user login even if the RADIUS server accepted the username and password.

To assign roles specifically to users on the RADIUS server, you configure the Accept Message for those accounts as either ips-role=administrator, ips-role=operator, ips-role=viewer, or ips-role=service. You configure the Accept Message individually for each user account. An example of a Reply attribute for a given user could be configured to return "Hello <user> your ips-role=operator."

If you configure a service account in the RADIUS server, you must also configure an identical service account locally on the device. For service accounts, both the RADIUS and Local accounts are checked during login.

Identifying an NTP Server

Use the NTP policy to configure a Network Time Protocol (NTP) server as the time source for the IPS device. Using NTP helps ensure synchronized time among your network devices, which can aid event analysis. NTP is the recommended way to configure time settings on an IPS device.

For detailed information on how to set the time on a sensor, including how to set up a Cisco IOS router as an NTP server, refer to Configuring Time in *Configuring the Cisco Intrusion Prevention System Sensor Using the Command Line Interface Version 7.0*.

Tip Check the time on your IPS sensor if you are having trouble updating your IPS software. If the time on the sensor is ahead of the time on the associated certificate, the certificate is rejected, and the sensor software update fails.

Step 1 Do one of the following to open the NTP policy:

- (Device view) Select **Platform > Device Admin > Server Access > NTP** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > NTP**, then select an existing policy or create a new one.

Step 2 In the **NTP Server IP Address** field, enter the IP address of the NTP server. You can also enter the name of a network/host object that identifies the single host address of the server, or click **Select** to select the object from a list or to create a new one.

Step 3 If the NTP server does not require authentication, deselect the **Authenticated NTP** checkbox.

If the NTP server requires authentication, configure the following options:

- **Authenticated NTP**—Select this option to enable authenticated connections.
- **Key, Confirm**—The key value of the NTP server. The key is an MD5 type of key (either numeric or character); it is the key that was used to set up the NTP server.

- **Key ID**—The key ID value of the NTP server, a numeric value between 1 and 65535.

Tip The key and key ID are configured on the NTP server; you must obtain them from the NTP server configuration.

Identifying DNS Servers

If you configure global correlation on an IPS 7.0+ sensor, the sensor must be able to resolve domain names to successfully connect to the update server when downloading global correlation updates. Use the DNS policy to identify the Domain Name System (DNS) servers that the sensor can use to resolve domain names to IP addresses.

Tip If your network requires HTTP proxies when making Internet connections, configure the HTTP Proxy policy instead of the DNS policy. See Identifying an HTTP Proxy Server, page 36-24.

Note The AIP-SSC-5 service module does not support DNS servers.

Step 1 Do one of the following to open the HTTP Proxy policy:

- (Device view) Select **Platform > Device Admin > Server Access > DNS** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > DNS**, then select an existing policy or create a new one.

Step 2 Specify the IP addresses of up to three DNS servers in the **Primary, Secondary, and Tertiary Address** fields. The sensor uses the servers in the order listed; if one server does not respond, the next server is contacted.

You can enter an IP address or the name of a network/host object that contains a server address. Click **Select** to select a network/host object from a list or to create a new one. The network/host object must specify a single host address.

Identifying an HTTP Proxy Server

If you configure global correlation on an IPS 7.0+ sensor, and your network requires the use of HTTP proxies to connect to the Internet, you need to configure the HTTP Proxy policy to identify a proxy that the IPS sensor can use. When downloading global correlation updates, the IPS sensor connects to the update server using this proxy. The proxy must be able to resolve DNS names.

Tip If you do not use HTTP proxies, configure DNS servers so that the IPS sensor can resolve the address of the update server. See Identifying DNS Servers, page 36-24.

Note The AIP-SSC-5 service module does not support HTTP proxy servers.

Step 1 Do one of the following to open the HTTP Proxy policy:

- (Device view) Select **Platform > Device Admin > Server Access > HTTP Proxy** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > HTTP Proxy**, then select an existing policy or create a new one.

Step 2 Configure the following options:

- **Enable Proxy**—Select this option to tell the device to connect through the configured proxy server.
- **IP Address**—Enter the IP address of the proxy server, or the name of the network/host object that contains the server's IP address. Click **Select** to select a network/host object from a list or to create a new one. The network/host object must contain a single host IP address.
- **Port**—Enter the port number used for HTTP connections to the proxy server. The default is 80.

IPS SSHv2 Known Host Keys

The IPS SSHv2 Known Host Keys policy enables you to configure SSHv2 server host keys (outgoing SSHv2 connections from an IPS sensor to an SSH server). This feature is available on IPS sensors running 7.1(8) and later versions of Cisco IPS.

The host key can be retrieved from an IPS sensor using valid IP addresses; alternatively, it can be entered manually if you know it. The host key retrieval may take few seconds.

Step 1 Do one of the following to open the IPS SSHv2 Known Host Keys policy:

- (Device view) Select **Platform > Device Admin > Server Access > SSHv2 Known Host Keys** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > SSHv2 Known Host Keys**, then select an existing policy or create a new one.

Step 2 Click the **Add** button to open the Add or Edit Known Host RSA Key Dialog Box, page 36-25.

Step 3 Select a row and then click the **Edit** button to open the Add or Edit Known Host RSA Key Dialog Box, page 36-25.

Add or Edit Known Host RSA Key Dialog Box

Use the Add or Edit Known Host RSA Key dialog box to retrieve an SSHv2 key from an IPS sensor or to enter the key manually if you know it.

Navigation Path

From the SSHv2 Known Host Keys policy, click the **Add** button beneath the IP Address/Public Key table, or select a row in the table and click the **Edit** button. For information on the SSHv2Known Host Keys policy, see IPS SSHv2 Known Host Keys, page 36-25.

Field Reference

Table 36-8 Add or Edit Known Host RSA Key Dialog Box

Element	Description
IP Address	The IP address of the IPS sensor from which you want to retrieve the public key.
Retrieve Public Key	Initiates retrieval of the public key from the device identified in the IP Address field. The Retrieve Public Key option is available in device view (it is not shown in shared policy view). However, you can enter inline values for the public key in shared policies or retrieve the public key in device view and share it using the "share policy" option.
Public Key	The public key that you know and are able to enter manually. For shared policies, you will be able to enter inline values for the host key.

Configuring IPS SSHv1 Fallback Settings

The IPS SSHv1 Fallback policy is available on IPS sensors running 7.1(8) and later versions of Cisco IPS.

Step 1 Do one of the following to enable or disable SSHv1 Fallback:

- (Device view) Select **Platform > Device Admin > Server Access >Settings** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > Settings**, then select an existing policy or create a new one.

Step 2 To enable SSHv1 fallback, click the checkbox.

Step 3 To disable SSHv1 fallback, clear the checkbox.

Configuring the External Product Interface

Use the External Product Interface policy to configure the way that Security Manager works with Management Center for Cisco Security Agents (CSA MC).

In general, the external product interface is designed to receive and process information from external security and management products. These external security and management products collect information that can be used to automatically enhance the sensor configuration information.

Management Center for Cisco Security Agents is the only external product that can be configured to communicate with the IPS. At most two Management Center for Cisco Security Agents servers can be configured per IPS device.

Tip Management Center for Cisco Security Agents is no longer an active product. Configure this policy only if you are still using that application. For more information, see About CSA MC in *Installing and Using Cisco Intrusion Prevention System Device Manager 6.0* and http://www.cisco.com/en/US/products/sw/cscowork/ps5212/index.html.

Management Center for Cisco Security Agents enforces a security policy on network hosts. It has two components:

- Agents that reside on and protect network hosts.
- A management console, which is an application that manages agents. It downloads security policy updates to agents and uploads operational information from agents.

Before You Begin

Add the external product as an allowed host so that Security Manager allows the sensor to communicate with the external product. For more information, see Identifying Allowed Hosts, page 36-7.

Step 1 Do one of the following to open the External Product Interface policy:

- (Device view) Select **Platform > Device Admin > Server Access > External Product Interface** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Server Access > External Product Interface**, then select an existing policy or create a new one.

The **Management Center for Cisco Security Agents** tab shows any existing definitions, including the IP address (or network/host object), URL, and port of the external application, the username and password used to log into it, and whether the connection is enabled. The interface type is always Extended SDEE.

Step 2 Do one of the following:

- To add a server, click the **Add Row (+)** button. This opens the External Product Interface dialog box. Enter the information required to identify the server and configure the posture ACLs. For detailed information on the settings, see External Product Interface Dialog Box, page 36-27.

  You can add at most two servers.
- To edit a server, select it and click the **Edit Row (pencil)** button and make the required changes in the External Product Interface dialog box.
- To delete a server, select it and click the **Delete Row (trash can)** button.

External Product Interface Dialog Box

Use the Add or Edit External Product Interface dialog box to add or modify interfaces between Management Center for Cisco Security Agents (CSA MC) and the IPS device and the related posture ACLs.

Navigation Path

From the External Product Interface IPS platform policy, click **Add Row** or select an entry and click **Edit Row**. For information on opening the External Product Interface policy, see Configuring the External Product Interface, page 36-26.

Field Reference

Table 36-9 External Product Interface Dialog Box

Element	Description
External Product's IP Address	The IP address, or the network/host policy object that contains the address, of the external product. Enter the IP address or object name, or click **Select** to select an object from a list or to create a new one.
Interface Type	Identifies the physical interface type, which is always Extended SDEE.
Enable receipt of information	Whether information is allowed to be passed from the external product to the sensor.
SDEE URL	The URL on the CSA MC the IPS uses to retrieve information using SDEE communication. You must configure the URL based on the software version of the CSA MC that the IPS is communicating with as follows: • For CSA MC version 5.0—/csamc50/sdee-server. • For CSA MC version 5.1—/csamc51/sdee-server. • For CSA MC version 5.2 and higher—/csamc/sdee-server (the default value).
Port	The port, or the port list object that identifies the port, being used for communications. Enter the port or port list name, or click **Select** to select the object from a list or to create a new object.
User name Password	A username and password that can log into the external product.
Enable receipt of host postures	Whether to allow the receipt of host posture information from CSA MC. The host posture information received from a CSA MC is deleted if you disable this option.
Allow unreachable hosts' postures	Whether to allow the receipt of host posture information for hosts that are not reachable by the CSA MC. A host is not reachable if the CSA MC cannot establish a connection with the host on any IP addresses in the host's posture. This option is useful in filtering the postures whose IP addresses may not be visible to the IPS sensor or that might be duplicated across the network. This filter is most applicable in network topologies where hosts that are not reachable by the CSA MC are also not reachable by the IPS, for example if the IPS and CSA MC are on the same network segment.

Table 36-9 External Product Interface Dialog Box (continued)

Element	Description
Posture ACL table	Posture ACLs are network addresses for which host postures are allowed or denied. Use posture ACLs to filter postures that have IP addresses that might not be visible to the IPS or that might be duplicated across the network. • To add a posture ACL, click the **Add Row (+)** button. This opens the Add Posture ACL dialog box. For information on configuring the Posture ACL, see Posture ACL Dialog Box, page 36-29. • To edit a posture ACL, select it and click the **Edit Row (pencil)** button. • To delete a posture ACL, select it and click the **Delete Row (trash can)** button. • To change the priority of an ACL, select it and click the Up or Down button. ACLs are processed in order, and the action associated with the first match is applied.
Enable receipt of watch listed addresses	Whether to allow the receipt of the watch list information from CSA MC. The watch list information received from a CSA MC is deleted if you disable this option.
Manual Watch List RR increase	The percentage of the manual watch list risk rating (RR). The default is 25, and the valid range is 0 to 35.
Session-based Watch List RR Increase	The percentage of the session-based watch list risk rating. The default is 25, and the valid range is 0 to 35.
Packed-based Watch List RR Increase	The percentage of the packet-based watch list risk rating. The default is 10, and the valid range is 0 to 35.

Posture ACL Dialog Box

Use the Add or Modify Posture ACL dialog box to configure posture ACLs for Management Center for Security Agents. Posture ACLs are network addresses for which host postures are allowed or denied. Use posture ACLs to filter postures that have IP addresses that might not be visible to the IPS or that might be duplicated across the network.

Configure the following fields to define a posture ACL:

- **Network Address**—Enter the IP address of a host or network, or the name of a network/host object that specifies one. You can click **Select** to select the object from a list or to create a new object.
- **Action**—Whether host postures will be permitted or denied from the hosts on the network address.

Navigation Path

From the External Product Interface dialog box (see External Product Interface Dialog Box, page 36-27), click the **Add Row (+)** button underneath the Posture ACL table, or select a posture ACL and click the **Edit Row (pencil)** button.

Configuring IPS Logging Policies

Use the IPS platform Logging policy to configure traffic flow notifications and Analysis Engine global variables. These settings apply to the general operation of the IPS sensor.

Traffic flow notifications have to do with the flow of traffic across the interface of a sensor. You can configure the sensor to monitor the flow of packets across an interface and send notification if that flow changes (starts and stops) during a specified interval. You can configure the missed packet threshold within a specific notification interval and also configure the interface idle delay before a status event is reported.

The Analysis Engine performs packet analysis and alert detection. It monitors traffic that flows through specified interfaces. For the Analysis Engine, there is only one global variable: Maximum Open IP Log Files.

Navigation Path

- (Device view) Select **Platform > Logging** from the Policy selector.
- (Policy view) Select **IPS > Platform > Logging**, then select an existing policy or create a new one.

Field Reference

Table 36-10 IPS Logging Page

Element	Description
Interface Notifications Tab	
Missed Packets Threshold	The percent of missed packets that has to occur before you want to receive notification. The default is 0, and the range is 0 to 100.
Notification Interval	The length of time, in seconds, that you want to check for the percentage of missed packets. The default is 30, and the range is 5 to 3600.
Interface Idle Threshold	The length of time, in seconds, that you will allow an interface to be idle and not receiving packets before you want to be notified. The default is 30, and the range is 5 to 3600.
Analysis Engine Tab	
Specify-Flow-Depth	Lets you specify the inspection depth of the flow. Flow depth is the number of bytes inspected in a flow. The new value applies for new flows only. The valid range is from 0 to 429496296. The default is 0.
Enable Service Activity	Service activity lets you gather information about service activities for diagnostic purposes. The details are more granular and have port level details. Enabling service activity impacts system performance. Enable service activity collection temporarily for diagnostic purposes only. You must reboot the sensor after you enable service activity for the change to take effect.
Service Activity Limit	Sets the limit for how many services you want to enable. The valid range is from 10 to 65536. The default is 15.
Note The Specify-Flow-Depth, Enable Service Activity, and Service Activity Limit fields are applicable to IPS devices of version 7.2(2) or higher.	
Maximum Open IP Log Files	The maximum number of open IP log files that you want to allow on the sensor. The default is 20, and the range is 20 to 100.

IPS Health Monitor

Use the IPS Health Monitor page to configure the metrics, or parameters, that are used to determine the health and network security status of your IPS devices. Your IPS devices use these metrics to assign appropriate severity when sending IPS events. The results appear in the Health and Performance Monitor of Security Manager (Launch > Health and Performance Monitor).

IPS Health Monitor is supported in IPS devices beginning with IPS version 6.1 and in Security Manager beginning with version 4.4. Please note the following special cases:

1. For IPS devices running 7.x, all 11 configuration items in the IPS Security Settings Policy are displayed and monitored properly in the Security Manager GUI.
2. For IPS devices running less than 6.1, the Network Participation and Global Correlation entries are hidden in the device view of Security Manager.
3. Some IPS Health Monitor configuration items are protected entries on the device side itself and cannot be edited. Security Manager informs you in such cases.

If you do not select a metric by checking the check box, it does not appear in the Health and Performance Monitor. You can accept the default configuration or edit the values. Items will be disabled and will not be editable if you do not select a metric.

The overall health is set to the most critical settings of any of the metrics. For instance, if all the selected metrics are normal except for one that is critical, the overall health becomes critical. The IPS sensor sends a health and security status event when the overall health status of the IPS sensor changes.

The security status of the IPS sensor is determined for each virtual sensor using the threat ratings of events detected by the virtual sensors. The security status of the virtual sensor is raised when the virtual sensor detects an event with a threat rating that exceeds the threshold for that virtual sensor. After a threshold has been exceeded, the security status remains at a critical level until the configured amount of time has passed with no more events being detected at the higher level.

To configure the metrics on the IPS Health Monitor page, select one of the following policies:

- (Device view) Select **Platform > Device Admin > Health Monitor** from the Policy selector.
- (Policy view) Select **IPS > Platform > Device Admin > Health Monitor** from the Policy Type selector, then select an existing policy or create a new one.

Note In policy view, no validation is performed if a shared IPS Health Monitor policy is applied to an IPS device running less than 6.1. Security Manager ignores such policies during deployment to device and captures them in deployment logs also.

The following table explains the IPS Health Monitor Metrics that you can configure.

Table 36-11 IPS Security Settings Policy

Element	Description
Inspection Load	Lets you set a threshold for inspection load and whether this metric is applied to the overall sensor health rating.
Missed Packet	Lets you set a threshold percentage for missed packets and whether this metric is applied to the overall sensor health rating.
Memory Usage	Lets you set a threshold percentage for memory usage and whether this metric is applied to the overall sensor health rating.

Table 36-11 IPS Security Settings Policy (continued)

Element	Description
Signature Update	Lets you set a threshold for when the last signature update was applied and whether this metric is applied to the overall sensor health rating.
License Expiration	Lets you set a threshold for when the license expires and whether this metric is applied to the overall sensor health rating.
Event Retrieval	Lets you set a threshold for when the last event was retrieved and whether this metric is applied to the overall sensor health rating. **Note** The event retrieval metric keeps track of when the last event was retrieved by an external monitoring application such as IME. Disable Event Retrieval if you are not doing external event monitoring.
Network Participation	Lets you choose whether the Network Participation health metrics contribute to the overall sensor health rating.
Global Correlation	Let you choose whether the Global Correlation health metrics contribute to the overall sensor health rating.
Application Failure	Lets you choose to have an application failure applied to the overall sensor health rating.
IPS in Bypass Mode	Let you choose to know if bypass mode is active and have that apply to the overall sensor health rating.
One or More Active Interfaces Down	Lets you choose to know if one or more enabled interfaces are down and have that apply to the overall sensor health rating.
Warning	Lets you set the lowest threshold in percentage, days, seconds, or failures for the warning threshold.
Critical	Lets you set the lowest threshold in percentage, days, seconds, or failures for the critical threshold.

Configuring IPS Security Settings

Use the IPS Security Settings policy to configure two items that are important to the security of your IPS devices:

- **Permit packet capture logging**—With this feature, IPS devices can prevent users from arbitrarily executing packet capture/display/iplog commands. In previous versions of Security Manager, such actions leave no trace of who executed the command.
- **Configurable idle timeout**—When configured, this feature terminates the connection to an IPS device after a period of time that you specify. Its purpose is to increase the security of a CLI session.

Note These settings are available for devices operating with IPS 7.1.3 and later.

To configure IPS security settings, select one of the following policies:

- (Device view) Select **Platform > Security >Settings** from the Policy selector.
- (Policy view) Select **IPS > Platform > Security > Settings** from the Policy Type selector, then select an existing policy or create a new one.

The following table explains the IPS security settings that you can configure.

Table 36-12 IPS Security Settings Policy

Element	Description
Permit packet logging	Whether to enable packet logging; applies to packet capture/display/iplog commands.
CLI Inactivity Timeout (In Minutes)	Terminates the connection to an IPS device after the specified period of time.

Managing IPS Device Interfaces

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

Dedicated IPS appliances and service modules have their own interface configuration, whereas Cisco IOS IPS devices are configured using the regular router interface policies. This chapter explains how to configure interfaces for dedicated IPS appliances and service modules only.

This chapter contains the following topics:

- Understanding Interfaces, page 37-1
- Understanding Interface Modes, page 37-2
- Configuring Interfaces, page 37-6

Understanding Interfaces

Tip This topic is an overview of IPS interfaces. For more detailed information, including the specific interface names and locations for each type of appliance and service module, supported roles, configuration restrictions, and hardware considerations, refer to the "Configuring Interfaces" chapter of the *Installing and Using Cisco Intrusion Prevention System Device Manager* for the IPS software version you are using on Cisco.com. The information is also in the IME and CLI guides. For general information, see http://www.cisco.com/go/ips.

The sensor interfaces are named according to the maximum speed and physical location of the interface. For example, GigabitEthernet2/1 supports a maximum speed of 1 Gigabit and is the second-from-the-right interface in the second-from-the bottom expansion slot.

There are three interface roles:

- **Command and control**—The command and control interface has an IP address and is used for configuring the sensor. It receives security and status events from the sensor and queries the sensor for statistics.

  The command and control interface is permanently enabled. It is permanently mapped to a specific physical interface, which depends on the specific model of sensor. You cannot use the command and control interface as either a sensing or alternate TCP reset interface. See the IPS document cited above for a list of command and control interfaces by device type.

- **Sensing**—Sensing interfaces are used by the sensor to analyze traffic for security violations. A sensor has one or more sensing interfaces depending on the sensor. Sensing interfaces can operate individually in promiscuous mode or you can pair them to create inline interfaces. In promiscuous mode, packets do not flow through the sensor; the sensor analyzes a copy of the monitored traffic. In inline mode, the IPS is in the traffic flow and can directly affect the traffic. For more information about sensing modes, see Understanding Interface Modes, page 37-2.

  Note On appliances, all sensing interfaces are disabled by default. You must enable them to use them. On modules, the sensing interfaces are permanently enabled. See the IPS document cited above for a list of sensing interfaces by device type.

- **Alternate TCP reset**—You can configure sensors to send TCP reset packets to try to reset a network connection between an attacker host and its intended target host. In some installations when the interface is operating in promiscuous mode, the sensor may not be able to send the TCP reset packets over the same sensing interface on which the attack was detected. In such cases, you can associate the sensing interface with an alternate TCP reset interface and any TCP resets that would otherwise be sent on the sensing interface when it is operating in promiscuous mode are instead sent out on the associated alternate TCP reset interface.

  If a sensing interface is associated with an alternate TCP reset interface, that association applies when the sensor is configured for promiscuous mode but is ignored when the sensing interface is configured for inline mode (interface or VLAN pair), because TCP resets are always sent on the sensing interfaces in those modes.

  Note With the exception of IDSM-2, any sensing interface can serve as the alternate TCP reset interface for another sensing interface. The alternate TCP reset interface on IDSM-2 is fixed because of hardware limitation. However, there is only one sensing interface on IPS modules (on routers or ASA devices), so you cannot specify an alternate TCP reset interface on them. See the IPS document cited above for a list of eligible alternate TCP reset interfaces by device type, and for more information about the conditions under which you would use one.

Understanding Interface Modes

Sensing interfaces can operate in various modes. The mode configured for an interface determines the traffic it can inspect and how it can respond to events.

This section contains the following topics:

- Promiscuous Mode, page 37-2
- Inline Interface Mode, page 37-3
- Inline VLAN Pair Mode, page 37-3
- VLAN Group Mode, page 37-4

Promiscuous Mode

In promiscuous mode, packets do not flow through the sensor. The sensor analyzes a copy of the monitored traffic rather than the actual forwarded packet. The advantage of operating in promiscuous mode is that the sensor does not affect the packet flow with the forwarded traffic. The disadvantage of

operating in promiscuous mode, however, is the sensor cannot stop malicious traffic from reaching its intended target for certain types of attacks, such as atomic attacks (single-packet attacks). The response actions implemented by promiscuous sensor devices are post-event responses and often require assistance from other networking devices, for example, routers and firewalls, to respond to an attack. While such response actions can prevent some classes of attacks, in atomic attacks the single packet has the chance of reaching the target system before the promiscuous-based sensor can apply an ACL modification on a managed device (such as a firewall, switch, or router).

By default, all sensing interfaces are in promiscuous mode. To change an interface from inline interface mode to promiscuous mode, delete any inline interface that contains that interface and delete any inline VLAN pair subinterfaces of that interface from the interface configuration.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Physical Interfaces, page 37-10

Inline Interface Mode

Operating in inline interface pair mode puts the IPS directly into the traffic flow and affects packet-forwarding rates making them slower by adding latency. This allows the sensor to stop attacks by dropping malicious traffic before it reaches the intended target, thus providing a protective service. Not only is the inline device processing information on Layers 3 and 4, but it is also analyzing the contents and payload of the packets for more sophisticated embedded attacks (Layers 3 to 7). This deeper analysis lets the system identify and stop or block attacks that would normally pass through a traditional firewall device.

In inline interface pair mode, a packet comes in through the first interface of the pair on the sensor and out the second interface of the pair. The packet is sent to the second interface of the pair unless that packet is being denied or modified by a signature.

Notes:

- If the paired interfaces are connected to the same switch, you should configure them on the switch as access ports with different access VLANs for the two ports. Otherwise, traffic does not flow through the inline interface.
- You can configure IPS modules for routers and ASA devices to operate inline even though these modules have only one sensing interface.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Inline Interface Pairs, page 37-13

Inline VLAN Pair Mode

You can associate VLANs in pairs on a physical interface. This is known as inline VLAN pair mode. Packets received on one of the paired VLANs are analyzed and then forwarded to the other VLAN in the pair.

Inline VLAN pair mode is an active sensing mode where a sensing interface acts as an 802.1q trunk port, and the sensor performs VLAN bridging between pairs of VLANs on the trunk. The sensor inspects the traffic it receives on each VLAN in each pair, and can either forward the packets on the other VLAN in

the pair, or drop the packet if an intrusion attempt is detected. You can configure an IPS sensor to simultaneously bridge up to 255 VLAN pairs on each sensing interface. The sensor replaces the VLAN ID field in the 802.1q header of each received packet with the ID of the egress VLAN on which the sensor forwards the packet. The sensor drops all packets received on any VLANs that are not assigned to inline VLAN pairs.

Notes:

- You cannot use the default VLAN as one of the paired VLANs in an inline VLAN pair.
- Inline VLAN pairs are not supported on IPS modules for routers or ASA devices.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Inline VLAN Pairs, page 37-14

VLAN Group Mode

You can divide each physical interface or inline interface into VLAN group subinterfaces, each of which consists of a group of VLANs on that interface. If you configure multiple virtual sensors, each of them can monitor one or more of these interfaces. This lets you apply multiple policies to the same sensor. The advantage is that now you can use a sensor with only a few interfaces as if it had many interfaces.

Note You cannot divide physical interfaces that are in inline VLAN pairs into VLAN groups.

VLAN group subinterfaces associate a set of VLANs with a physical or inline interface. No VLAN can be a member of more than one VLAN group subinterface. Each VLAN group subinterface is identified by a number between 1 and 255. Subinterface 0 is a reserved subinterface number used to represent the entire unvirtualized physical or logical interface. You cannot create, delete, or modify subinterface 0 and no statistics are reported for it.

When you create a VLAN group, it is either promiscuous or inline:

- Promiscuous VLAN group—If you configure a VLAN group on a physical interface, the VLAN group is promiscuous, as described in Promiscuous Mode, page 37-2.
- Inline VLAN group—If you configure a VLAN group on an inline interface pair (a logical interface), the VLAN group is inline, as described in Inline Interface Mode, page 37-3.

Thus, VLAN groups augment the operation of promiscuous mode interfaces or inline interfaces by confining their operation to selected VLANs. Once you assign a VLAN group to an interface (physical or inline interface), the interface is no longer a plain promiscuous or inline interface pair and can only be used for inline VLAN groups.

An unassigned VLAN group is maintained that contains all VLANs that are not specifically assigned to another VLAN group. You cannot directly specify the VLANs that are in the unassigned group. When a VLAN is added to or deleted from another VLAN group subinterface, the unassigned group is updated.

Packets in the native VLAN of an 802.1q trunk do not normally have 802.1q encapsulation headers to identify the VLAN number to which the packets belong. A default VLAN variable is associated with each physical interface and you should set this variable to the VLAN number of the native VLAN or to 0. The value 0 indicates that the native VLAN is either unknown or you do not care if it is specified. If the default VLAN setting is 0, the following occurs:

- Any alerts triggered by packets without 802.1q encapsulation have a VLAN value of 0 reported in the alert.
- Non-802.1q encapsulated traffic is associated with the unassigned VLAN group and it is not possible to assign the native VLAN to any other VLAN group.

Note You can configure a port on a switch as either an access port or a trunk port. On an access port, all traffic is in a single VLAN is called the access VLAN. On a trunk port, multiple VLANs can be carried over the port, and each packet has a special header attached called the 802.1q header that contains the VLAN ID. This header is commonly referred as the VLAN tag. However, a trunk port has a special VLAN called the native VLAN. Packets in the native VLAN do not have the 802.1q headers attached. IDSM-2 can read the 802.1q headers for all nonnative traffic to determine the VLAN ID for that packet. However, IDSM-2 does not know which VLAN is configured as the native VLAN for the port in the switch configuration, so it does not know what VLAN the native packets are in. Therefore, you must tell IDSM-2 which VLAN is the native VLAN for that port. Then IDSM-2 treats any untagged packets as if they were tagged with the native VLAN ID.

Related Topics

- Deploying VLAN Groups, page 37-5
- Understanding Interfaces, page 37-1
- Configuring VLAN Groups, page 37-15

Deploying VLAN Groups

Because a VLAN group of an inline pair does not translate the VLAN ID, an inline paired interface must exist between two switches to use VLAN groups on a logical interface. For an appliance, you can connect the two pairs to the same switch, make them access ports, and then set the access VLANs for the two ports differently. In this configuration, the sensor connects between two VLANs, because each of the two ports is in access mode and carries only one VLAN. In this case the two ports must be in different VLANs, and the sensor bridges the two VLANs, monitoring any traffic that flows between the two VLANs. IDSM-2 also operates in this manner, because its two data ports are always connected to the same switch.

You can also connect appliances between two switches. There are two variations. In the first variation, the two ports are configured as access ports, so they carry a single VLAN. In this way, the sensor bridges a single VLAN between the two switches.

In the second variation, the two ports are configured as trunk ports, so they can carry multiple VLANs. In this configuration, the sensor bridges multiple VLANs between the two switches. Because multiple VLANs are carried over the inline interface pair, the VLANs can be divided into groups and each group can be assigned to a virtual sensor. The second variation does not apply to IDSM-2 because it cannot be connected in this way.

Related Topics

- Understanding Interfaces, page 37-1
- VLAN Group Mode, page 37-4
- Configuring VLAN Groups, page 37-15

Configuring Interfaces

Use the Interfaces policy for IPS appliances and service modules to configure the interface settings for the device. The following topics explain how to configure the various types of settings. These topics do not apply to Cisco IOS IPS devices, which use the standard router interface policies.

- Understanding the IPS Interfaces Policy, page 37-6
- Configuring Physical Interfaces, page 37-10
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Inline Interface Pairs, page 37-13
- Configuring Inline VLAN Pairs, page 37-14
- Configuring VLAN Groups, page 37-15
- Viewing a Summary of IPS Interface Configuration, page 37-8

Understanding the IPS Interfaces Policy

Use the Interfaces policy to configure the physical interfaces, inline pairs, VLAN pairs, and VLAN groups on IPS appliances and service modules. This policy does not apply to Cisco IOS IPS devices.

You can configure any single physical interface to run in promiscuous mode, inline pair mode, inline VLAN pair mode, promiscuous VLAN group, or inline VLAN group, but you cannot configure an interface in a combination of these modes.

Tip The contents of this policy differ depending on the device type and IPS software version. For example, some devices display the physical interfaces tab only; creating the other types of configurations is not supported. If a tab or option described below does not appear on the policy you are configuring, it does not apply to the device.

Navigation Path

(Device view only) Select **IPS > Interfaces** from the Policy selector.

Related Topics

- Understanding Interfaces, page 37-1
- Understanding Interface Modes, page 37-2
- Discovering Policies on Devices Already in Security Manager, page 5-15

Field Reference

Table 37-1 IPS Interfaces Policy

Element	Description
Physical Interfaces tab	The physical interfaces that are available on the device. You can edit these interfaces only (select the device and click the **Edit Row** button); you must perform inventory discovery on the device to obtain the correct list of physical interfaces, for example, if you add an interface card to the device. The columns displayed on the tab show the configuration of each interface and are explained in Modify Physical Interface Map Dialog Box, page 37-11. Note that the Administrative State column indicates whether the interface is enabled (Yes or No); you must enable an interface for it to function. For more information, see Configuring Physical Interfaces, page 37-10.
Inline Pairs tab	The inline interface pairs that allow inline mode processing, as described in Inline Interface Mode, page 37-3. The table shows the name of the pair, the interfaces that are part it, and a description, if any. For more information, see Configuring Inline Interface Pairs, page 37-13. • To add a pair, click the **Add Row** button and fill in the Add Interface Pair dialog box. • To edit a pair, select it and click the **Edit Row** button. • To delete a pair, select it and click the **Delete Row** button.
VLAN Pairs tab	The VLAN pairs for each physical interface, as described in Inline VLAN Pair Mode, page 37-3. The table shows the interface and subinterface, with the two VLANs that are paired, and a description, if any. For more information, see Configuring Inline VLAN Pairs, page 37-14. • To add a pair, click the **Add Row** button and fill in the Add VLAN Pair dialog box. • To edit a pair, select it and click the **Edit Row** button. • To delete a pair, select it and click the **Delete Row** button.
VLAN Groups tab	The VLAN groups defined for a physical interface or inline pair, as described in VLAN Group Mode, page 37-4. The table shows the name of the interface or pair, the VLAN group (empty means all unassigned VLANs), and a description, if any. For more information, see Configuring VLAN Groups, page 37-15. • To add a group, click the **Add Row** button and fill in the Add VLAN Group dialog box. • To edit a group, select it and click the **Edit Row** button. • To delete a group, select it and click the **Delete Row** button.

Table 37-1 IPS Interfaces Policy (continued)

Element	Description
Summary tab	A summary of how you have configured the sensing interfaces—the interfaces you have configured for promiscuous mode, the interfaces you have configured as inline pairs, and the interfaces you have configured as inline VLAN pairs. For more information, see Viewing a Summary of IPS Interface Configuration, page 37-8.
Bypass Mode	The bypass mode for the device, which determines how the sensor should handle inline mode traffic when the sensor processes are temporarily stopped for upgrades or when the sensor monitoring processes fail. This is a global setting that applies to all inline mode interfaces on the device. Select the desired option; for a detailed explanation of how each of these options affect inline traffic, see Configuring Bypass Mode, page 37-12. • **Off (Always inspect inline traffic)**—Disables bypass mode. Traffic is always inspected, and if the monitoring process of the sensor is down, traffic stops flowing. • **On (Never inspect inline traffic)**—Traffic bypasses the Analysis Engine and is never inspected. • **Auto (Bypass inspection when analysis engine is stopped)**—Traffic is inspected unless the monitoring process of the sensor is down, in which case traffic continues to flow through the sensor uninspected. This is the default. Auto mode is useful during sensor upgrades to ensure that traffic is still flowing while the sensor is being upgraded.
CDP Mode	How to handle Cisco Discovery Protocol (CDP) packets. The CDP configuration applies globally to all interfaces on the device, however, it has an effect only on inline interfaces (both inline interfaces and inline VLAN pairs). For more information, see Configuring CDP Mode, page 37-13. Select the desired option: • **Forward CDP packets**—To allow CDP packets to pass through the sensor. • **Drop CDP packets**—To have the sensor drop all CDP packets and not allow them to pass through the sensor. This is the default setting.

Viewing a Summary of IPS Interface Configuration

The Summary tab of the Interfaces policy contains a table summarizing how you have configured the sensing interfaces—the interfaces you have configured for promiscuous mode, the interfaces you have configured as inline pairs, the interfaces you have configured as inline VLAN pairs, inline VLAN groups, and promiscuous VLAN groups. The content of this table changes when you change your interface configuration.

You can configure any single physical interface to run in promiscuous mode, inline pair mode, or inline VLAN pair mode, but you cannot configure an interface in a combination of these modes.

Tip Not all service modules have a summary tab.

Navigation Path

(Device view) Select **Interfaces** from the Policy selector. Click the **Summary** tab.

Related Topics

- Understanding Interfaces, page 37-1
- Understanding the IPS Interfaces Policy, page 37-6
- Configuring Physical Interfaces, page 37-10
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Inline Interface Pairs, page 37-13
- Configuring Inline VLAN Pairs, page 37-14
- Configuring VLAN Groups, page 37-15

Field Reference

Table 37-2 IPS Interface Summary Tab

Element	Description
Name	The name of the interface. The names are FastEthernet or GigabitEthernet for promiscuous interfaces. For inline interfaces, the name is whatever you assigned to the pair.
Subinterface Number	The subinterface number of the inline VLAN pair or VLAN group. Subinterface numbers can be from 1 to 255.
Inline Interface Name	The name of the inline interface pair.
Mode	The mode for the interface: promiscuous, inline, promiscuous VLAN group, or inline VLAN group and whether there are VLAN pairs. For an explanation of interface modes, see Understanding Interface Modes, page 37-2.
VLAN A VLAN B	The VLAN ID for the first and second VLANs for VLAN pairs. VLAN numbers can be from 1 to 4095.
VLAN Range	The range of VLAN IDs belonging to the VLAN group, for example, 100-200. If the VLAN group is configured to apply to all unassigned VLANs, the field is empty.

Configuring Physical Interfaces

The Physical Interfaces tab of the IPS Interfaces policy lists the existing physical interfaces on your sensor and their associated settings. You cannot add or delete physical interfaces in this policy; instead, you must use policy discovery to obtain the current list of interfaces from the device. Thus, if you add or remove interface cards (available for some appliances), you must rediscover the device as described in Discovering Policies on Devices Already in Security Manager, page 5-15.

To configure the sensor to monitor traffic, you must enable the interface using this procedure. When you initialized the sensor using the **setup** command (using the command line interface on the IPS), you assigned the interface or the inline pair to a virtual sensor, and enabled the interface or inline pair. If you need to change your interfaces settings, you can do so on the Physical Interfaces tab. To assign an interface to a virtual sensor, select the Virtual Sensors policy and add or edit the virtual sensor, as appropriate.

Tip Each physical interface can be divided into VLAN group subinterfaces, each of which consists of a group of VLANs on that interface. For more information, see Configuring VLAN Groups, page 37-15.

Related Topics

- Understanding Interfaces, page 37-1
- Defining A Virtual Sensor, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9
- Assigning Interfaces to Virtual Sensors, page 38-4
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Inline Interface Pairs, page 37-13

Step 1 (Device view) Select **Interfaces** from the Policy selector, then click the **Physical Interfaces** tab (if necessary).

Step 2 Select the interface whose configuration you want to change and click the **Edit Row** button. The Modify Physical Interface Map dialog box appears.

Step 3 Make the desired configuration changes and click **OK**. Following are the settings you are most likely to want to change; for a description of all options, see Modify Physical Interface Map Dialog Box, page 37-11.

- **Enabled**—Whether the interface is enabled (**Yes** or **No**). Select Yes to make the interface functional. The value of this option is shown in the Administrative State column in the Physical Interfaces tab.
- **Default VLAN**—The VLAN to which the interface is assigned.
- **Specify Interface for TCP Reset**—If you want to assign an alternate TCP reset interface, as described in Understanding Interfaces, page 37-1, select this option, then select the alternate interface from the **interface-name** list.

Modify Physical Interface Map Dialog Box

Use the Modify Physical Interface Map dialog box to change the configuration of the physical interfaces of an IPS sensor. For the procedure, see Configuring Physical Interfaces, page 37-10.

Navigation Path

(Device view) Select **Interfaces** from the Policy selector. On the **Physical Interfaces** tab, select an interface and click the **Edit Row** button.

Related Topics

- Understanding Interfaces, page 37-1
- Understanding the IPS Interfaces Policy, page 37-6

Field Reference

Table 37-3 Modify Physical Interface Map Dialog Box

Element	Description
Name	The name of the physical interface.
Media Type	The type of media for the physical interface. The media types are the following: • TX—Copper media. • SX—Fiber media. • XL—Network accelerator card. • Backplane interface—An internal interface that connects the module to the backplane of the parent chassis.
Description	A description of the interface.
Enabled	Whether the interface is enabled, Yes or No. You must select **Yes** for the interface to be functional. You also have to assign the interface to a virtual sensor for it to monitor traffic; use the Virtual Sensors policy.
Duplex	The duplex setting of the interface. The duplex types are the following: • Auto—Sets the interface to auto negotiate duplex. • Full—Sets the interface to full duplex. • Half—Sets the interface to half duplex.
Speed	The speed setting of the interface. The speed options are the following: • Auto—Sets the interface to auto negotiate speed. • 10 MB—Sets the interface to 10 MB (for TX interfaces only). • 100 MB—Sets the interface to 100 MB (for TX interfaces only). • 1 GB—Sets the interface to 1 GB (for gigabit interfaces only). • 10 GB—Sets the interface to 10 GB (for 10 gigabit interfaces only).
Default VLAN	The VLAN ID associated with native traffic, or 0 if unknown or if you do not care which VLAN it is.

Table 37-3 Modify Physical Interface Map Dialog Box (continued)

Element	Description
Specify Interface for TCP Reset interface-name	Whether to send TCP resets on an alternate interface when this interface is used for promiscuous monitoring and the reset action is triggered by a signature firing. If you select this option, select the alternate TCP reset interface from the interface-name list. For more information about alternate TCP reset, see Understanding Interfaces, page 37-1.

Configuring Bypass Mode

You can use inline bypass as a diagnostic tool and a failover protection mechanism. Normally, the sensor Analysis Engine performs packet analysis. When inline bypass is activated, Analysis Engine is bypassed, allowing traffic to flow through the inline interfaces and inline VLAN pairs without inspection. Inline bypass ensures that packets continue to flow through the sensor when the sensor processes are temporarily stopped for upgrades or when the sensor monitoring processes fail. There are three modes: on, off, and automatic. By default, bypass mode is set to automatic.

Keep the following factors in mind before deciding which bypass mode to use:

- There are security consequences when you put the sensor in bypass mode. When bypass mode is on, the traffic bypasses the sensor and is not inspected; therefore, the sensor cannot prevent malicious attacks.
- The inline bypass functionality is implemented in software, so it functions only when the operating system is running. If the sensor is powered off or shut down, inline bypass does not work—traffic does not flow through the sensor.
- When the sensor applies a signature or global correlation update, it might trigger bypass. Whether bypass is triggered depends on the traffic load of the sensor and the size of the signature or global correlation update. If bypass mode is turned off, an inline sensor stops passing traffic while the update is being applied.

To change the bypass mode setting, follow these steps:

Step 1 (Device view) Select the **Interfaces** policy from the Policy selector.

Step 2 In the **Bypass Mode** field at the bottom of the policy, select the desired option:

- **Off (Always inspect inline traffic)**—Disables bypass mode.

  Traffic flows through the sensor for inspection. If the monitoring process of the sensor is down, traffic stops flowing. This means that inline traffic is always inspected.
- **On (Never inspect inline traffic)**—Traffic bypasses the Analysis Engine and is not inspected. This means that inline traffic is never inspected.
- **Auto (Bypass inspection when analysis engine is stopped)**—Traffic flows through the sensor for inspection unless the monitoring process of the sensor is down. This is the default.

If the monitoring process of the sensor is down, traffic bypasses the sensor until the sensor is running again. The sensor then inspects the traffic. Auto mode is useful during sensor upgrades to ensure that traffic is still flowing while the sensor is being upgraded. Auto mode also helps to ensure traffic continues to pass through the sensor if the monitoring process fails.

Configuring CDP Mode

You can configure the IPS sensor to enable or disable the forwarding of Cisco Discovery Protocol (CDP) packets. The CDP configuration applies globally to all interfaces on the device, however, it has an effect only on inline interfaces (both inline interfaces and inline VLAN pairs).

Cisco Discovery Protocol is a media- and protocol-independent device-discovery protocol that runs on all Cisco-manufactured equipment, including routers, access servers, bridges, and switches. Using CDP, a device can advertise its existence to other devices and receive information about other devices on the same LAN or on the remote side of a WAN. CDP runs on all media that support SNAP, including LANs, Frame Relay, and ATM media.

Tip The CDP Mode setting is not available on all IPS appliances and service modules. If the CDP Mode field does not appear on the Interfaces policy, the setting does not apply to the device you are configuring.

To change the CDP mode setting on a device, follow these steps:

Step 1 (Device view) Select the **Interfaces** policy from the Policy selector.

Step 2 In the **CDP Mode** field at the bottom of the policy, select the desired option:

- **Forward CDP packets**—To allow CDP packets to pass through the sensor.
- **Drop CDP packets**—To have the sensor drop all CDP packets and not allow them to pass through the sensor. This is the default setting.

Configuring Inline Interface Pairs

You can pair interfaces on your sensor if your sensor is capable of inline monitoring. For more information about inline pairs, see Inline Interface Mode, page 37-3.

Tip IPS modules for routers and ASA devices do not need an inline pair for monitoring. You only need to add the physical interface to a virtual sensor.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Physical Interfaces, page 37-10
- Configuring VLAN Groups, page 37-15

- Defining A Virtual Sensor, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9
- Assigning Interfaces to Virtual Sensors, page 38-4

Step 1 (Device view) Select **Interfaces** from the Policy selector, then click the **Inline Pairs** tab.

Step 2 Do one of the following:

- To add a pair, click the **Add Row** button. The Add Interface Pair dialog box opens.
- To edit a pair, select it and click the **Edit Row** button. The Edit Interface Pair dialog box opens.

Tip You can also delete a pair by selecting it and clicking the **Delete Row** button. You cannot delete an inline pair if there is an inline VLAN group. First delete the inline VLAN group from the VLAN Groups tab, and then delete the inline pair.

Step 3 In the Add or Edit Inline Pairs dialog box, configure the following options:

- **Inline Interface Name**—The name you want to give to this inline pair. The name cannot be longer than 32 characters; alphanumeric and underscore characters are allowed. You cannot edit this name after you create the pair.
- **Interface 1 and 2**—Select the two physical interfaces that you want to form a pair. The lists include only those interfaces that are defined on the Physical Interfaces tab and that are not already part of an inline pair, VLAN pair, or VLAN group.
- **Description**—An optional description for the pair.

Step 4 Click **OK** to save your changes.

Configuring Inline VLAN Pairs

Use the VLAN Pairs tab of the IPS Interfaces policy to configure the VLAN pairs for physical interfaces. The summary table displays the existing VLAN pairs for each physical interface. You can create multiple VLAN pairs on a single physical interface. For more information about inline VLAN pair mode, see Inline VLAN Pair Mode, page 37-3.

Tips

- You cannot create a VLAN pair for an interface if it is already part of an inline interface pair; create VLAN groups for inline interface pairs.
- To create an inline VLAN pair for an interface that is in promiscuous mode and assigned to a virtual sensor, you must first remove the interface from the virtual sensor (using the Virtual Sensors policy) and then create the inline VLAN pair.
- You cannot use the default VLAN as one of the paired VLANs in an inline VLAN pair.
- If your sensor does not support inline VLAN pairs, the VLAN Pairs pane is not displayed. IPS modules on routers and ASA devices do not support inline VLAN pairs.
- When using inline VLAN pairs, you should configure UniDirectional Link Detection (UDLD) on the connected switch that is hosting the VLANs. UDLD can help switches prevent spanning-tree forwarding loops and single direction links. For detailed information, see Configuring UDLD in *Installing and Using Intrusion Prevention System Device Manager*.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Physical Interfaces, page 37-10
- Configuring VLAN Groups, page 37-15

Step 1 (Device view) Select **Interfaces** from the Policy selector, then click the **VLAN Pairs** tab.

Step 2 Do one of the following:

- To add a pair, click the **Add Row** button. The Add VLAN Pair dialog box opens.
- To edit a pair, select it and click the **Edit Row** button. The Edit VLAN Pair dialog box opens.

Tip You can also delete a pair by selecting it and clicking the **Delete Row** button. You cannot delete an inline VLAN pair if it is assigned to a virtual sensor. First remove the assignment to the virtual sensor using the Virtual Sensors policy, and then delete the inline VLAN pair.

Step 3 In the Add or Edit VLAN Pairs dialog box, configure the following options:

- **Physical Interfaces**—Select the physical interface on which you are creating this VLAN pair. The list includes only those interfaces that are defined on the Physical Interfaces tab and that are not already part of an inline interface pair or VLAN group. However, you can create multiple VLAN pairs on a single interface.
- **Subinterface Number**—Enter a number to assign as a subinterface. The number must be unique on the interface, that is, it cannot already be assigned to another VLAN pair on the selected physical interface. Subinterface numbers can be from 1 to 255.
- **Description**—An optional description for the pair.
- **VLAN A, B**—The numbers of the two VLANs that you want to join as a pair. VLAN numbers are from 1 to 4095. You must enter different numbers, and the numbers must not already be part of another VLAN pair on the selected physical interface.

Step 4 Click **OK** to save your changes.

Configuring VLAN Groups

Use the VLAN Groups tab of the IPS Interfaces policy to configure the VLAN groups for physical interfaces and inline interface pairs (logical interfaces). The summary table displays the existing VLAN groups. You can create multiple VLAN groups on a single physical interface or inline interface pair. For more information about VLAN group mode, see VLAN Group Mode, page 37-4.

A VLAN group consists of a group of VLAN IDs that exist on an interface. Each VLAN group consists of at least one VLAN ID. You can have up to 255 VLAN groups per interface (logical or physical). Each group can contain any number of VLAN IDs.

After you assign the VLAN IDs to the VLAN group, you must assign the VLAN group to a virtual sensor for it to be operational. You can assign a single group to at most one virtual sensor. Use the Virtual Sensors policy to make the assignment.

Note VLAN groups are supported in IPS 6.0 and later only. Not all IPS appliances or service modules support VLAN groups. If the VLAN Groups tab does not appear in the Interfaces policy, the device you are configuring does not support the feature.

Related Topics

- Understanding Interfaces, page 37-1
- Configuring Bypass Mode, page 37-12
- Configuring CDP Mode, page 37-13
- Configuring Physical Interfaces, page 37-10
- Defining A Virtual Sensor, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9
- Assigning Interfaces to Virtual Sensors, page 38-4

Step 1 (Device view) Select **Interfaces** from the Policy selector, then click the **VLAN Groups** tab.

The table shows the existing VLAN groups, including the interface for which the group is defined, the subinterface number, description (if any), and the VLANs assigned to the group. If the VLANs cell is empty, the group is defined for all unassigned VLANs on the interface.

Step 2 Do one of the following:

- To add a pair, click the **Add Row** button. The Add VLAN Group dialog box opens.
- To edit a pair, select it and click the **Edit Row** button. The Edit VLAN Group dialog box opens.

Tip You can also delete a group by selecting it and clicking the **Delete Row** button. You cannot delete a VLAN group if it is assigned to a virtual sensor. First remove the assignment to the virtual sensor using the Virtual Sensors policy, and then delete the VLAN group.

Step 3 In the Add or Edit VLAN Group dialog box, configure the following options:

- **Physical and Logical Interfaces**—Select the physical interface or inline interface pair for which you are creating this VLAN group. The list includes only unpaired physical interfaces (defined on the Physical Interfaces tab) that do not already have inline VLAN pairs defined, or inline interface pairs that are defined on the Inline Pairs tab. You can create multiple VLAN groups on a single interface. Keep the following in mind:
  - If you select a physical interface, you are creating a promiscuous VLAN group.
  - If you select a logical interface, you are creating an inline VLAN group.
- **Subinterface Number**—Enter a number to assign as a subinterface. The number must be unique on the interface, that is, it cannot already be assigned to another VLAN group on the selected interface. Subinterface numbers can be from 1 to 255.
- **Description**—An optional description for the group.
- VLAN assignment—Select one of the following options:
  - **All Unassigned VLAN IDs**—The group contains all VLANs that are not assigned to other VLAN groups. This is the default option

- **Range of free VLAN IDs**—The group contains specific VLANs. In the **Range** box, enter any combination of single VLAN IDs or ranges (separate starting and ending ID with a hyphen), and separate multiple entries with commas. For example, 10, 12-25, 33-49. VLAN numbers are from 1 to 4095.

  The VLAN ID cannot already be in another VLAN group for the selected interface. The VLANs also must be configured on the connected switch or there will be no traffic to inspect.

Step 4 Click **OK** to save your changes.

CHAPTER 38

Configuring Virtual Sensors

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

All IPS devices and service modules have a base virtual sensor named vs0. When you configure the IPS appliance or service module, you must configure the base vs0 sensor to assign interfaces to it. This assignment tells the device which interfaces to inspect. There are also other settings that are configured on virtual sensors.

In addition to the base vs0 virtual sensor, many IPS appliances and service modules allow you to create user-defined virtual sensors. You can use these virtual sensors to create separate policies for traffic, so that a single physical sensor can act as if it were multiple sensors. A virtual sensor is a logical grouping of sensing interfaces and the configuration policy for the signature engines and event action filters to apply.

This chapter contains the following topics:

- Understanding the Virtual Sensor, page 38-1
- Defining A Virtual Sensor, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9
- Deleting A Virtual Sensor, page 38-10

Understanding the Virtual Sensor

The sensor can receive data inputs from one or many monitored data streams. These monitored data streams can either be physical interface ports or virtual interface ports. For example, a single sensor can monitor traffic from in front of the firewall, from behind the firewall, or from in front of and behind the firewall concurrently. A single sensor can monitor one or more data streams. In this situation a single sensor policy or configuration is applied to all monitored data streams.

With virtual sensors, you can create separate policies to apply to specific traffic feeds. For example, if you want to create a policy for a data center and a second much different policy for the campus network, yet run both policies on the same hardware device, you can configure separate virtual sensors to implement these policies.

You configure the following policies and settings separately for a virtual sensor:

- Signature and signature settings (policies in the IPS > Signatures folder).
- Event action policies (policies in the IPS > Event Actions folder).

- Anomaly detection policies (the IPS > Anomaly Detection policy) and the anomaly detection mode (in the Virtual Sensors policy).
- The promiscuous interfaces, inline interface pairs, inline VLAN pairs, inline VLAN groups, or promiscuous VLAN groups that the virtual sensor monitors.

Note No packet is processed by more than one virtual sensor; you cannot assign the same physical or logical interface to more than one sensor. Packets from interfaces, inline interface pairs, inline VLAN pairs, and VLAN groups that are not assigned to any virtual sensor are disposed of according to the inline bypass configuration that you define in the **Interfaces** policy.

- The inline TCP session tracking and Normalizer modes (in the Virtual Sensors policy).

Note If you create a policy instance on an IPS device for signatures, event actions, or anomaly detection but do not assign it to any of the virtual sensors on that device (that is, you do not use that policy instance), then that policy instance is deleted by Security Manager during deployment.

All other policies and settings are configured on the parent device that hosts the virtual sensor. For example, if you want to use global correlation, you configure it on the parent device and the virtual sensors share that configuration.

You can configure up to four virtual sensors on one appliance, but you can add only three user-defined virtual sensors. The first virtual sensor, vs0, is the base sensor and you cannot delete it. In Security Manager, virtual sensors are presented as follows:

- The device selector in Device view contains the parent device, which doubles as the base virtual sensor, vs0. Select this device to configure all device-level policies and to create virtual sensors in the Virtual Sensors policy.
- Any user-defined virtual sensors are also shown in the device selector in Device view. The display name of the real device is prepended to the beginning of the name of the virtual sensor. In most cases, the result is that the virtual sensors appear next to the parent (real) device that the virtual sensor is on. For example, on the host (real device) named “bob,” the virtual sensor with the name “vs1” will appear in the device list as “bob_vs1.”

  To configure the signature, anomaly detection, and event action policies for a virtual sensor, you must select it in the device selector. You cannot configure these policies by selecting the parent device; those policies on the parent device are for the vs0 base sensor.

The following topics explain more about virtual sensors:

- Advantages and Restrictions of Virtualization, page 38-3
- Inline TCP Session Tracking Mode, page 38-3
- Understanding Normalizer Mode, page 38-4
- Assigning Interfaces to Virtual Sensors, page 38-4
- Identifying the Virtual Sensors for a Device, page 38-5
- Defining A Virtual Sensor, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9
- Deleting A Virtual Sensor, page 38-10

Advantages and Restrictions of Virtualization

An advantage of using virtual sensors is that you can operate more than one virtual sensor on one appliance while configuring each virtual sensor differently with regard to signature behavior and traffic feed. For example, if you want to create a policy for a data center and a second much different policy for the campus network, yet run both policies on the same hardware device, you can configure separate virtual sensors to implement these policies.

Virtualization has the following advantages:

- You can apply different configurations to different sets of traffic.
- You can monitor two networks with overlapping IP spaces with one sensor.
- You can monitor both inside and outside a firewall or NAT device.

Virtualization has the following restrictions:

- You must assign both sides of asymmetric traffic to the same virtual sensor.
- Using VACL capture or SPAN (promiscuous monitoring) is inconsistent with regard to VLAN tagging, which causes problems with VLAN groups.
  - When using Cisco IOS software, a VACL capture port or a SPAN target does not always receive tagged packets even if it is configured for trunking.
  - When using the MSFC, fast path switching of learned routes changes the behavior of VACL captures and SPAN.
- Persistent store is limited.
- Not all IPS sensors support multiple virtual sensors. The Virtual Sensors policy appears for all IPS appliances and service modules, because you must use it to assign interfaces to the base vs0 sensor. If the Add button in the policy is disabled for a device, and you have not configured user-defined virtual sensors, then the device does not support virtualization. Examples of devices that do not support virtualization include the Cisco IPS 4215, NM-CIDS, AIM-IPS, NME-IPS, and AIP-SSC. IDSM2 supports virtualization, but it does not support VLAN groups or inline interface pairs.
- You must use IPS 6.0+ software. Older software versions do not support virtualization.
- Cisco IOS IPS devices do not support virtualization. Use the **IPS > Interface Rules** policy to specify the interfaces that IPS should monitor.

Virtualization has the following traffic capture requirements:

- The virtual sensor must receive traffic that has 802.1q headers (other than traffic on the native VLAN of the capture port).
- The sensor must see both directions of traffic in the same VLAN group in the same virtual sensor for any given sensor.

Related Topics

- Understanding the Virtual Sensor, page 38-1
- Defining A Virtual Sensor, page 38-5

Inline TCP Session Tracking Mode

When you choose to modify packets inline, if the packets from a stream are seen twice by the Normalizer engine, it cannot properly track the stream state and often the stream is dropped. This situation occurs most often when a stream is routed through multiple VLANs or interfaces that are being monitored by

the IPS. A further complication in this situation is the necessity of allowing asymmetric traffic to merge for proper tracking of streams when the traffic for either direction is received from different VLANs or interfaces.

To deal with this situation, you can set the mode so that streams are perceived as unique if they are received on separate interfaces or VLANs (or the subinterface for VLAN pairs).

The following inline TCP session tracking modes apply:

- **Interface and VLAN**—All packets with the same session key (AaBb) in the same VLAN (or inline VLAN pair) and on the same interface belong to the same session. Packets with the same key but on different VLANs are tracked separately.
- **VLAN Only**—All packets with the same session key (AaBb) in the same VLAN (or inline VLAN pair) regardless of the interface belong to the same session. Packets with the same key but on different VLANs are tracked separately.
- **Virtual Sensor**—All packets with the same session key (AaBb) within a virtual sensor belong to the same session. This is the default and almost always the best option to choose.

You configure the inline TCP session tracking mode as a property of the virtual sensor as described in Defining A Virtual Sensor, page 38-5.

Understanding Normalizer Mode

Normalizer mode applies only when the sensor is operating in inline mode. The default is strict evasion protection, which is full enforcement of TCP state and sequence tracking. The Normalizer enforces duplicate packets, changed packets, out-of-order packets, and so forth, which helps prevent attackers from evading the IPS.

Asymmetric mode disables most of the Normalizer checks. Use Asymmetric mode only when the entire stream cannot be inspected, because in this situation, attackers can now evade the IPS.

You configure the Normalizer mode as a property of the virtual sensor as described in Defining A Virtual Sensor, page 38-5.

Assigning Interfaces to Virtual Sensors

An IPS sensor monitors traffic that traverses interfaces, interface pairs, or VLAN pairs assigned to a virtual sensor.

You can assign one or more of the following types of interfaces to a virtual sensor:

- Promiscuous interface—A physical interface that does not have VLAN groups and which is not part of an inline interface pair.
- Inline interface pair—A logical interface composed of two physical interfaces.
- Inline VLAN pair—A logical interface composed of two VLANs.
- Promiscuous VLAN group—A VLAN group that is assigned to a subinterface on a physical interface.

  The physical interface cannot already be used for an inline interface or VLAN pair. There can be many promiscuous VLAN groups on the same promiscuous interface, but the VLANs assigned cannot overlap. Once a VLAN group is assigned to a promiscuous interface, it is no longer a plain promiscuous interface and can only be used for promiscuous VLAN groups.
- Inline VLAN group—A VLAN group that is assigned to a subinterface of an inline interface pair.

There can be many inline VLAN groups on the same inline interface pair, but the VLANs assigned cannot overlap. Once a VLAN group is assigned to an inline interface pair it is no longer a plain inline interface pair and can only be used for inline VLAN groups.

VLAN groups cannot be assigned to inline VLAN pairs.

You must configure the interfaces before you can assign them to virtual sensors. For more information about configuring all of these types of interfaces, see Configuring Interfaces, page 37-6. For information on assigning interfaces to virtual sensors, see Defining A Virtual Sensor, page 38-5.

Identifying the Virtual Sensors for a Device

If you configure user-defined virtual sensors on an IPS appliance or service module, the virtual sensor appears in the device selector in Device view.

Normally, the display name of a virtual sensor is in the form *device-name_virtual-sensor-name*, where *device-name* is the name of the parent device, and *virtual-sensor-name* is the name of the virtual sensor. For example, the virtual sensor vs1 on device 10.100.10.10 would be 10.100.10.10_vs1.

Thus, under normal conditions, the virtual sensors for a device should appear immediately after the parent device in the device selector. However, you can change the virtual sensor's display name by editing the device properties. If you alter the default name, the virtual sensors might not appear anywhere near the parent device in the device selector.

You can use the following techniques to identify the virtual sensors defined on a device, or to identify the parent device of a virtual sensor:

- To see a list of virtual sensors defined on an IPS device, select the **Virtual Sensors** policy on the device. The table shows all virtual sensors, including the base vs0 sensor. Note that the vs0 sensor does not appear separately in the device selector; it is represented by the parent device itself.

  Unless you radically alter the display names of virtual sensors, the virtual sensor name, along with the parent device's display name, should help you find the virtual sensor in the device selector.
- To determine which IPS device is the host of a virtual sensor, right-click the virtual sensor in the device selector and select **Device Properties**. The Hostname display-only field on the General tab shows the host device display name plus the virtual sensor name as defined on the device.

Defining A Virtual Sensor

Use the Virtual Sensors policy to configure virtual sensors on your Cisco IPS devices. Even if your IPS device does not support multiple virtual sensors, you must use this policy to assign interfaces to the base sensor, vs0, and configure properties that are associated with the virtual sensor.

Tip For Cisco IOS IPS devices, you configure the interfaces that the IPS examines in the **IPS > Interface Rules** policy. You cannot configure virtual sensors in an IOS IPS device.

Before You Begin

Configure the interfaces on the sensor, including inline interface pairs, inline VLAN pairs, and promiscuous and inline VLAN groups. The interface configurations must exist before you can assign them to a virtual sensor. For information on interfaces, interface modes, and how to configure them, see Chapter 37, "Managing IPS Device Interfaces".

Related Topics

- Understanding Interfaces, page 37-1
- Understanding Interface Modes, page 37-2
- Advantages and Restrictions of Virtualization, page 38-3
- Inline TCP Session Tracking Mode, page 38-3
- Understanding Normalizer Mode, page 38-4
- Assigning Interfaces to Virtual Sensors, page 38-4
- Identifying the Virtual Sensors for a Device, page 38-5
- Editing Policies for a Virtual Sensor, page 38-9

Step 1 (Device view only.) Select **Virtual Sensors** from the Policies selector to open the Virtual Sensors policy.

The policy lists all existing virtual sensors, including the base vs0 sensor, which you cannot delete. The information for each sensor shows the interfaces assigned to the sensor, anomaly detection mode, inline TCP tracking mode, normalizer mode, and a description, if any. If the Assignments cell is empty, no interfaces are assigned to the virtual sensor, which means the virtual sensor cannot analyze any traffic.

Step 2 Do one of the following:

- To add a virtual sensor, click the **Add Row** button. The Add Virtual Sensor dialog box opens.

  You can add at most three sensors. The device supports four virtual sensors, including the base vs0 sensor. If the Add Row button is disabled, you either have configured the maximum number of sensors, or your device does not support multiple virtual sensors.
- To edit a virtual sensor, select it and click the **Edit Row** button. The Edit Virtual Sensor dialog box opens.

Tip You can also delete a virtual sensor by selecting it and clicking the **Delete Row** button. You cannot delete the base vs0 sensor. For more information about deleting virtual sensors, see Deleting A Virtual Sensor, page 38-10.

Step 3 In the Add or Edit Virtual Sensor dialog box, configure at least the following options. The defaults for the other options are appropriate in most cases. For detailed information on all available options, see Virtual Sensor Dialog Box, page 38-7.

- Virtual Sensor Name—The name of the virtual sensor. The virtual sensor name can be up to 64 characters and it cannot contain spaces.
- Interface assignments (Available, Assigned lists)—The promiscuous interfaces, inline interface pairs, inline VLAN pairs, promiscuous VLAN groups, or inline VLAN groups that you want this virtual sensor to use. The list of available interfaces shows only those interfaces that are configured in the Interfaces policy and that are not yet assigned to another virtual sensor.
  - To assign interfaces, select them in the available list and click **>>**.
  - To remove an assignment, select the interface in the assigned list and click **<<**. You must remove an assignment before you can assign an interface to a different virtual sensor.

  Tip: If you are not sure about the content of a specific interface, for example, its mode or assigned VLANs, close the dialog box, go to the Interfaces policy, and examine the various tabs.

Step 4 Click **OK** to save your changes and add them to the Virtual Sensors policy.

Step 5 Click **Save** to save the Virtual Sensors policy.

Step 6 If you created a new virtual sensor, you must submit your changes to the database for the new virtual sensor to appear in the device selector in Device view.

- Non-Workflow mode—Select **File > Submit**.
- Workflow mode—Select **Activities > Approve Activity**, or if you are operating with an activity approver, **Activities > Submit Activity**. The activity must be approved before the virtual sensor appears in the device selector.

Note In the device selector, the display name of the real device is prepended to the beginning of the name of the virtual sensor. In most cases, the result is that the virtual sensors appear next to the parent (real) device that the virtual sensor is on. For example, on the host (real device) named "bob," the virtual sensor with the name "vs1" will appear in the device list as "bob_vs1."

Step 7 To configure the policies associated with a virtual sensor, select it in the device selector in Device view. You can then configure the associated policies. See the following topics:

- Chapter 39, "Defining IPS Signatures"
- Chapter 40, "Configuring Event Action Rules"
- Configuring Anomaly Detection, page 41-6

Virtual Sensor Dialog Box

Use the Add or Edit Virtual Sensor dialog box to configure the properties for a virtual sensor.

Navigation Path

(Device view only.) Select **Virtual Sensors** from the Policy selector. Click the **Add Row** button, or select an existing virtual sensor and click the **Edit Row** button.

Related Topics

- Defining A Virtual Sensor, page 38-5
- Advantages and Restrictions of Virtualization, page 38-3
- Assigning Interfaces to Virtual Sensors, page 38-4
- Chapter 37, "Managing IPS Device Interfaces"
- Understanding Interfaces, page 37-1
- Understanding Interface Modes, page 37-2

Field Reference

Table 38-1 Add or Edit Virtual Sensor Dialog Box

Element	Description
Virtual Sensor Name	The name of the virtual sensor. The virtual sensor name can be up to 64 characters and it cannot contain spaces. The name of the default virtual sensor is **vs0**. You cannot change the name after you create the virtual sensor. To change a virtual sensor name, delete the sensor and create a new sensor with the desired name. If you already configured local policies for the sensor (that is, signature, event action, and anomaly detection policies), first save the policies as shared policies, delete the sensor, create the new sensor, then assign the shared policies to the new virtual sensor. For more information about creating shared policies from local policies, see Sharing a Local Policy, page 5-41.
Interface Assignments (Available, Assigned)	The promiscuous interfaces, inline interface pairs, inline VLAN pairs, promiscuous VLAN groups, or inline VLAN groups that you want this virtual sensor to use. The list of available interfaces shows only those interfaces that are configured in the Interfaces policy and that are not yet assigned to another virtual sensor. • To assign interfaces, select them in the available list and click **>>**. • To remove an assignment, select the interface in the assigned list and click **<<**. You must remove an assignment before you can assign an interface to a different virtual sensor. **Tip** If you are not sure about the content of a specific interface, for example, its mode or assigned VLANs, close the dialog box, go to the Interfaces policy, and examine the various tabs.
Anomaly Detection Mode	The mode that you want the anomaly detection policy to operate in for this virtual sensor: Detect, Inactive, Learn. The default and normal operational mode is Detect. However, if you are using asymmetric normalizer mode, you might want to set the anomaly detection mode to Inactive. For detailed information about these modes, see Anomaly Detection Modes, page 41-2.
Inline TCP Session Tracking Mode	The mode used to segregate multiple views of the same stream if the same stream passes through the sensor more than once. The default mode is Virtual Sensor. For more information, see Inline TCP Session Tracking Mode, page 38-3. Select one of the following: • Interface and VLAN—All packets with the same session key (AaBb) in the same VLAN (or inline VLAN pair) and on the same interface belong to the same session. Packets with the same key but on different VLANs are tracked separately. • VLAN Only—All packets with the same session key (AaBb) in the same VLAN (or inline VLAN pair) regardless of the interface belong to the same session. Packets with the same key but on different VLANs are tracked separately. • Virtual Sensor—All packets with the same session key (AaBb) within a virtual sensor belong to the same session.

Table 38-1 Add or Edit Virtual Sensor Dialog Box (continued)

Element	Description
Normalizer Mode	The type of Normalizer mode you need for traffic inspection. For more information, see Understanding Normalizer Mode, page 38-4. • Strict Evasion Protection—(Default) If a packet is missed for any reason, all packets after the missed packet are not processed. Strict evasion protection provides full enforcement of TCP state and sequence tracking. Any out-of-order packets or missed packets can produce Normalizer engine signatures 1300 or 1330 firings, which try to correct the situation, but can result in denied connections. • Asymmetric Mode Protection—Can only see one direction of bidirectional traffic flow. Asymmetric mode protection relaxes the evasion protection at the TCP layer. Asymmetric mode lets the sensor synchronize state with the flow and maintain inspection for those engines that do not require both directions. Asymmetric mode lowers security because full protection requires both sides of traffic to be seen.
Description	The description of the virtual sensor.

Editing Policies for a Virtual Sensor

Virtual sensors have two types of policies: the virtual sensor's properties, and policies assigned to the virtual sensor. You use a different approach to edit these items.

- To edit the properties of a virtual sensor, select the virtual sensor's parent device in the device selector in Device view. Then, select the **Virtual Sensors** policy. You can then select the virtual sensor in the table and click the **Edit Row** button.

  Using the Virtual Sensors policy, you can change the interfaces assigned to a sensor, the anomaly detection mode, inline TCP session tracking mode, and Normalizer mode. For more information, see the following topics:

  – Defining A Virtual Sensor, page 38-5
  – Virtual Sensor Dialog Box, page 38-7

- To edit the policies assigned to a virtual sensor, select the virtual sensor in the device selector in Device view. A virtual sensor's name is in the form *device-name_virtual-sensor-name*, where *device-name* is the name of the parent device, and *virtual-sensor-name* is the name of the virtual sensor. For example, the virtual sensor vs1 on device 10.100.10.10 would be 10.100.10.10_vs1.

  Note The base virtual sensor, vs0, is integrated with the parent device and does not appear separately in the device selector. To configure the base virtual sensor, select the parent device.

  You can then select the policies in the Policies selector and configure them. For more information, see the following topics:

  – Chapter 39, "Defining IPS Signatures"

- Chapter 40, "Configuring Event Action Rules"
- Configuring Anomaly Detection, page 41-6

All other policies are configured on the parent device, and the configurations apply to all virtual sensors configured on the device.

Deleting A Virtual Sensor

Virtual sensors appear in the device selector in Device view. However, you cannot delete them from the selector using the same command used for other devices. Instead, you must delete the virtual sensor from the Virtual Sensors policy of the parent device, that is, the device on which the virtual sensor is defined. The following procedure explains how to delete a user-defined virtual sensor.

Tip The base virtual sensor, vs0, does not appear in the device selector. Instead, it is represented by the parent IPS sensor; it is considered to be the base IPS device. To delete the base vs0 sensor, you delete the entire device from the inventory. For information on deleting devices from the inventory, see Deleting Devices from the Security Manager Inventory, page 3-58.

Before You Begin

When you delete a virtual sensor, you also delete the policies defined for the sensor, such as signature, event action, and anomaly detection policies. If you configured non-default local policies, and you want to preserve them for use on other virtual sensors, you must first convert the local policies to shared policies. Then, after you delete the virtual sensor, the policies continue to exist as unassigned shared policies. You can then assign them to another virtual sensor. For more information on creating a shared policy from a local policy, see Sharing a Local Policy, page 5-41.

Using this technique is ideal if you are deleting a virtual sensor simply as a means to change the virtual sensor's name. Because you cannot change a virtual sensor's name, you must delete it and create a new virtual sensor with the desired name. If you created shared policies, you could then assign those shared policies to your new sensor, and it will have the same configuration as the sensor had under the old name.

Step 1 (Device view only.) Select **Virtual Sensors** from the Policies selector to open the Virtual Sensors policy.

Step 2 Select the user-defined virtual sensor that you want to delete and click the **Delete Row** button.

Step 3 You are asked for a two-step confirmation. First, you are warned that you must save the policy to keep the policy and device synchronized. Click **OK** to continue, and you are also asked to confirm that you want to delete the node.

If you confirm, the virtual sensor is removed from both the policy and the device selector. It will take a few moments before the device view is updated and the virtual sensor disappears from the list of devices.

Defining IPS Signatures

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

You can use Security Manager to configure IPS signatures for dedicated IPS appliances and service modules or Cisco IOS IPS devices. When configuring signatures for Cisco IOS IPS, keep in mind that the router cannot use as many signatures as a dedicated appliance or service module.

This chapter contains the following topics:

- Understanding Signatures, page 39-1
- Configuring Signatures, page 39-4
- Configuring Signature Settings, page 39-30

Understanding Signatures

Network intrusions are attacks on, or other misuses of, network resources. Cisco IPS sensors and Cisco IOS IPS devices use a signature-based technology to detect network intrusions. A signature specifies the types of network intrusions that you want the sensor to detect and report. As sensors scan network packets, they use signatures to detect known types of attacks, such as denial of service (DoS) attacks, and respond with actions that you define.

On a basic level, signature-based intrusion detection technology can be compared to virus-checking programs. Cisco IPS contains a set of signatures that the sensor compares with network activity. When a match is found, the sensor takes some action, such as logging the event or sending an alarm to the Security Manager Event Viewer.

Signatures can produce false positives, because certain normal network activity can be construed as malicious. For example, some network applications or operating systems may send out numerous ICMP messages, which a signature-based detection system might interpret as an attempt by an attacker to map out a network segment. You can minimize false positives by editing your signature parameters (tuning your signatures).

To configure a sensor to monitor network traffic for a particular signature, you must enable the signature. By default, the most critical signatures are enabled when you install the signature update. When an attack is detected that matches an enabled signature, the sensor generates an alert, which is stored in the event store of the sensor. The alerts, as well as other events, may be retrieved from the event store by web-based clients such as Event Viewer. By default the sensor logs all Informational alerts or higher.

Some signatures have subsignatures, that is, the signature is divided into subcategories. When you configure a subsignature, changes made to the parameters of one subsignature apply only to that subsignature. For example, if you edit signature 3050 subsignature 1 and change the severity, the severity change applies to only subsignature 1 and not to 3050 2, 3050 3, and 3050 4.

Cisco IPS contains over 10,000 built-in default signatures. You cannot rename or delete signatures from the list of built-in signatures, but you can retire signatures to remove them from the sensing engine. You can later activate retired signatures; however, this process requires the sensing engines to rebuild their configuration, which takes time and could delay the processing of traffic. You can tune built-in signatures by adjusting several signature parameters. Built-in signatures that have been modified are called tuned signatures.

Note We recommend that you retire any signatures that you are not using. This improves sensor performance.

You can create signatures, which are called custom signatures. Custom signature IDs begin at 60000. You can configure them for several things, such as matching of strings on UDP connections, tracking of network floods, and scans. Each signature is created using a signature engine specifically designed for the type of traffic being monitored.

For more about signatures, see:

- Obtaining Detailed Information About a Signature, page 39-2
- Understanding Signature Inheritance, page 39-3

Related Topics

- Configuring Signatures, page 39-4
- Chapter 42, “Configuring Global Correlation”

Obtaining Detailed Information About a Signature

You can find detailed information about each signature from the Cisco Security Intelligence Operations web site. The web site includes a wealth of information and best practice recommendations for network security, and you can set up IntelliShield alerts. There is education on advanced security topics to help you protect your network, prioritize remediation, and structure your systems to reduce organizational risk.

When you edit the Signatures policy in Security Manager (see Signatures Page, page 39-4), the signature ID is linked directly into the Cisco Security Intelligence Operations database of IPS signatures. Clicking a signature ID opens a page containing information about the signature, including a description, the vulnerabilities on which the signature is based, when the signature was created, and so forth. You can search this database yourself at http://tools.cisco.com/security/center/search.x?search=Signature. (The database was formerly called the Cisco Network Security Database or NSDB.)

If you do not have access to Cisco.com, then the signature ID is linked to a local copy of the signature database information. Security Manager detects whether you have access to Cisco.com and makes the appropriate link for you without your having to set a preference.

The database includes information only for built-in, default signatures. You cannot find information about custom (user-defined) signatures.

Beginning with Security Manager 4.4, the Signatures Page (IPS > Signatures > Signatures) contains an Explanation tab and a Related Threats tab for each signature. These tabs display detailed information in a separate window on the Signatures page. For example, the Explanation tab displays Description, Signature ID, and so forth; the Related Threats tab displays vulnerabilities for other software that you may be using, and so forth.

Tip If this window is not visible to you, expand it with the up arrow button in the bottom-left corner of the Signatures page. To hide this window, collapse it with the corresponding down arrow, also in the bottom-left corner of the Signatures page. You can resize this window with standard controls.

Understanding Signature Inheritance

Signature inheritance for IPS devices is different than for any other Security Manager rules-based policy. Inheritance refers to the capability of Security Manager to enforce hierarchical lists of first-match, rule-based policies such as access rules. Signature inheritance is different because for IPS devices, Security Manager allows inheritance on a per-signature basis.

This example shows what is meant by inheritance on a per-signature basis:

Step 1 In Policy View, select **IPS > Signatures > Signatures**.

Step 2 Create a policy named test1.

Step 3 Create a second policy, named test2.

Step 4 Right-click **test 2** and select **Inherit Signatures**. The Inherit Rules—test 2 dialog box appears.

Step 5 Select **test1** and click the **OK** button.

Step 6 Select **test1** and edit a signature. Note the edit that you made and save your change.

Step 7 Select **test2** and select the signature that you just edited. Observe that test2 inherited the editing that you did on test1.

IPS Signature Purge

Beginning with Security Manager 4.1, old signature versions (defined as being older than the lowest signature level deployed) are purged during a periodic purge operation, the purpose of which is to optimize the database.

Note As a result of the purge operation, you may notice the deletion of some of your unused tuning contexts.

Some of the purged signatures may be restored during your next download of IPS signature packages from Cisco.com.

IPS signature purge is disabled by default. To enable IPS signature purge,

Step 1 Stop the Cisco Security Manager Daemon Manager: At the command prompt, enter **net stop crmdmgtd**.

Step 2 Navigate to *NMSROOT*\MDC\ips\etc\sensorupdate.properties file, where *NMSROOT* is the path to the Security Manager installation directory. The default is C:\Program Files\CSCOpx.

Step 3 In sensorupdate.properties, change purgeUnusedSignautesEntriesinDB:false to purgeUnusedSignautesEntriesinDB:true.

Step 4 Re-start the Cisco Security Manager Daemon Manager: At the command prompt, enter **net start crmdmgtd**.

IPS signature purge now runs at midnight every day.

Configuring Signatures

The Signatures policy is where you configure signatures for Cisco IPS sensors and Cisco IOS IPS devices.

This section contains the following topics:

- Signatures Page, page 39-4
- Viewing Signature Update Levels, page 39-13
- Enabling and Disabling Signatures, page 39-14
- Editing Signatures, page 39-14
- Adding Custom Signatures, page 39-19
- Cloning Signatures, page 39-21
- Regular Expressions in Custom Signatures, page 39-22
- Editing Signature Parameters (Tuning Signatures), page 39-23

Signatures Page

Use the Signatures page to display the signature summary table, in which you can add, edit, and delete IPS signatures. From this page you can tune the active signature set in a policy by enabling or disabling signatures. You can also use this page to unload signatures from the engine.

Beginning with Version 4.6, Security Manager enables you to apply a signature threat profile to one or more signature policies, starting from IPS device version 7.3(1). A signature threat profile is a predefined signature template that includes customized tunings. These tunings adjust the signature coverage and response actions to enable the sensor to make better choices in various deployment and threat scenarios. This Signatures page displays the threat profile and its version, that has been applied to the policy. Click the **To Change** button to select a threat profile to apply to the policy. For more information, see Apply Signature Threat Profiles, page 39-9. To see the signatures that belong to a threat profile, filter the Source column to contain the text **Threat Profile**. For information about how to filter tables, see Filtering Tables, page 1-48

If you download a particular signature package that does not contain one or more of the threat profiles already created in shared signature policy, Security Manager displays the warning message "**Currently applied threat profile is not applicable to this signature version**" on the shared signature Policy View. Similarly on the Device View, Security Manager displays the same warning message if you try to apply the shared signature policy to an unsupported device.

Since threat profile updates cannot be performed separately, you must update the current signature version of the device if you want to update its threat profile version. Note that any update made to a threat profile version modifies the signatures associated with the threat profiles, but retains any user-defined signature tuning already performed by the user.

Note Threat profiles are not supported in IOS-IPS.

Tips

- Enabled and disabled signatures are indicated by the "Enabled" checkbox for a particular signature. In previous releases of Security Manager, disabled signatures were indicated by hash marks covering the table row. When you deploy the configuration, disabled signatures are removed from the device. For more information, see Enabling and Disabling Signatures, page 39-14.
- For many columns, you can right-click the column and edit the property directly. Your edits apply to all rows that you have selected. If you select more than one row, the options that you can select are limited to those that are valid for all selected rows. The contents of the right-click menu differ on the basis of the cell that you right-click. For more information on the available commands, see Signature Shortcut Menu, page 39-10.
- To show or hide a column, right-click the table heading row on the signature summary table and then click **Show Columns**. By default, all columns are shown.

  Note Beginning with Version 4.5, Security Manager has a Notes column for each signature; this feature enables you to add a note so that you can revisit particular signatures later to see what you or other users have added for a signature or an event. This feature is helpful for network administrators in monitoring noisy signatures or signatures that need particular attention. However, the Notes column may not appear by default after you restore a Security Manager database. To show the Notes column, right-click the table heading row on the signature summary table, then click **Show Columns**, and finally click **Notes**. You may discover this situation during installation of Security Manager if you back up and restore the database; however, this situation will not occur during inline upgrades.

Navigation Path

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

Related Topics

- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49
- Understanding Signatures, page 39-1
- Understanding Signature Inheritance, page 39-3
- Enabling and Disabling Signatures, page 39-14
- Cloning Signatures, page 39-21
- Editing Signature Parameters (Tuning Signatures), page 39-23
- Configuring Event Action Filters, page 40-4
- Chapter 40, “Configuring Event Action Rules”

Field Reference

Table 39-1 Signature Policy

Element	Description
ID	The signature ID, which is the unique numerical value assigned to this signature. This value lets the sensor identify a particular signature. Click the ID number to open a page in your web browser with detailed information about the signature, as explained in Obtaining Detailed Information About a Signature, page 39-2.
Sub	The subsignature ID, which is the unique numerical value assigned to this subsignature. A Subsignature ID identifies a more granular version of a broad signature.
Name	The name assigned to the signature.
Enabled	A checkbox indicating whether the signature is enabled or disabled in this policy. A signature must be enabled for the sensor to protect against the traffic specified by the signature.
Severity	The severity level that the signature reports: High, Medium, Low, or Informational.
Fidelity	The weight associated with how well this signature might perform in the absence of specific knowledge of the target.

Table 39-1 Signature Policy (continued)

Element	Description
Notes	Enables you to add a note so that you can revisit particular signatures later to see what you or other users have added for a signature or an event. This feature is helpful for network administrators in monitoring noisy signatures or signatures that need particular attention. Notes are not saved to the device during deployment to the device. Notes as described here are a Security Manager GUI feature only and are disregarded during deployment to the device, as they are not part of any IPS policy in Security Manager. Notes are not part of any IPS policy, so assignment or inheritance of a shared signature policy will have no effect on Notes. Right-clicking a signature and adding notes will not prompt for activity/ticket creation. However, adding notes by double-clicking a signature or clicking the Edit button will prompt for activity/ticket creation because this involves signature policy modification. Notes cannot be added to signatures as part of a signature update operation. Other parameters can be edited, though. Notes cannot be searched by the Global Search feature. The Notes column will not display the Notes text. Only an icon will be displayed. You must double-click the icon to display the Notes text. If you use the "Export to File" button, then in the resultant .csv file, the Notes column will display only Y or N, and in this way signifies that those signatures were annotated. The actual text will not be exported. Notes cannot be edited. All added notes are appended to the existing notes as a new note entry. You can of course delete the note and add the updated note afresh. To add a note, right-click the row for a particular signature and then click **Add Note**. After you add the note, click **Save** and then close the Notes dialog box. After you close the Notes dialog box, the row for a particular signature will display a "Note" icon. To add a note to more than one signature, select the signatures you want (Shift-click or Ctrl-click in Windows) and proceed as you would for one particular signature. Notes can be local or shared. If you add a note to a device with only local policies, you can add, edit, and delete only local notes. If you add a note to a device with a shared policy assigned, you can add, edit, and delete both local and shared notes—check the **Share this Note** option. However, you can add notes only to *that particular device* (i.e., local override of notes) even if a shared policy is assigned. **Tip** If you have a shared policy assigned to a device and want to add the notes only to that device for a particular signature without affecting the shared policy, then you need to add the notes without choosing the **Share this Note** option in Device View. **Tip** You can also work with notes in the Edit Signature dialog box. Refer to Edit Button later in this table.

Table 39-1 Signature Policy (continued)

Element	Description
Base RR	The base risk rating value of the signature.
Actions	The actions the sensor takes when this signature fires.
Source	The lowest policy in the inheritance hierarchy that overrides the settings for a signature. Values can be: • Default—The signature uses the default Cisco-defined settings. • Local—The signature is defined specifically for the selected device (Device view only.) • Policy name—The lowest shared policy in the inheritance hierarchy. You can see policy names in Policy view, or in Device view if you assign a shared signature policy to the device.
Retired	The conditions under which the signature is retired, if any. A retired signature is removed from the signature engine. You can activate a retired signature to place it back in the signature engine. **Timesaver** Use the retired field to unload disabled signatures on your IOS-IPS device to achieve the most favorable memory consumption of that device. If the engine level of a signature policy is less than E-4, the Retired field has two possible values: false and true. False means that the signature is not retired; true means that the signature is retired. If the engine level of a signature policy is equal to E-4, the Retired field has four possible values: • **false**—The signature is not retired. • **low-mem-retired**—The signature should be retired on low-memory platforms. A low-memory device is one that has 2 GB RAM or less. • **med-mem-retired**—The signature should be retired on both low-end and medium platforms. A medium-memory device is one that has 4 GB RAM or less, but more than 2 GB RAM. (Any device with more than 4 GB RAM is considered a high-memory platform.) • **true**—The signature is retired on all platforms. When you select low-mem-retired or med-mem-retired, Security Manager configures the device with those signatures. Whether the signature is actually retired on the device depends on amount of memory installed on the device; the device makes the decision on which signatures are actually retired. **Tip** The term *engine level* used here is not the same as the term *engine* in the row above.
Engine	The engine that parses and inspects the traffic specified by this signature.
View Update Level button (Device view only.)	Click this button view the signature update level for this device. For more information, see Viewing Signature Update Levels, page 39-13.

Table 39-1 Signature Policy (continued)

Element	Description
Export to File button	Click this button to export the signature summary for the current device to a comma-separated values (CSV) file. You are prompted to select the folder on the Security Manager server and to specify a file name.
Add button	Click this button to add a custom signature. For more information, see the following topics: • Adding Custom Signatures, page 39-19 • Edit Signature or Add Custom Signature Dialog Boxes, page 39-15
Edit button	Click this button to edit the selected signature. You can edit one signature at a time. For more information, see the following topics: • Editing Signatures, page 39-14 • Edit Signature or Add Custom Signature Dialog Boxes, page 39-15
Delete button	Click this button to delete the selected custom signatures. You cannot delete Cisco-defined signatures. If you do not want to deploy a Cisco-defined signature, you can retire it or disable it.

Apply Signature Threat Profiles

Use the Apply Threat Profile dialog box to select a signature threat profile from the available profiles and apply to the policy. Applying a threat profile modifies only the **Enabled** and **Retired** fields on the Signatures Page, page 39-4. After you have applied a particular threat profile to a policy, the corresponding signature tunings are merged with the existing signatures on the Signatures page. To see the signatures that belong to a threat profile, on the Signatures page, filter the Source column to contain the text **Threat Profile**. For information about how to filter tables, see Filtering Tables, page 1-48

Select any of the following threat profiles that are currently provided by Cisco:

- SCADA—Select this threat profile template if you are using the Cisco IPS device primarily to protect industrial control systems. In addition to signatures in the default set, SCADA signature template includes specialized signatures for general SCADA protocol detections and specific identifiers that address tools and environments common to most device controlled environments.
- Edge—Select this threat profile template if you are using the Cisco IPS device primarily for securing an internet connection. In addition to signatures in the default set, Edge signature template includes additional signatures that provide broader protection for desktop operating systems, web browsers, web technologies, and common desktop applications.
- Web_Applications—Select this threat profile template if you are using the Cisco IPS device primarily for protecting web server farms. In addition to signatures in the default set, Web_Applications signature template includes additional signatures that provide broader protection for web servers, web development tools and frameworks, content management systems, load balancers, and databases.
- Data Center—Select this threat profile template if you are using the Cisco IPS device primarily for protecting data centers. In addition to signatures in the default set, Data Center signature template includes additional signatures that provide broader protection for server operating systems, web servers, application servers, databases, content management systems, messaging servers and virtualization systems.

Note Any signature tuning performed by the user on Local signatures (signatures defined for a selected device and for which the source policy is Local) will be preserved over the threat profile. For Default signatures, the threat profile tunings will be preserved.

Navigation Path

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS devices) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.

Field Reference

Table 39-2 Threat Profile Details

Element	Description
Signature ID	The signature ID, which is the unique numerical value assigned to this signature. This value lets the sensor identify a particular signature.
Sub Signature ID	The sub signature ID, which is the unique numerical value assigned to this sub signature. A sub signature ID identifies a more granular version of a broad signature.
Enabled	Indicates whether the signature is enabled or disabled in this threat profile. A signature must be enabled for the sensor to protect against the traffic specified by the signature.
Retired	Indicates whether the signature is retired or active in this threat profile.
Has Conflict	Indicates whether a signature with tunings from the applied threat profile also has tunings performed by the user. Signatures that are tuned by the user and by the applied threat profile are flagged as True in the Has Conflict column. If there is no conflict between the applied threat profile and user tunings for a signature, the Has Conflict column for that signature shows False. **Note** Any signature tuning performed by the user on Local signatures (signatures defined for a selected device and for which the source policy is Local) will be preserved over the threat profile. For Default signatures, the threat profile tunings will be preserved.

Signature Shortcut Menu

Right-clicking inside the signature summary table in the Signatures policy displays a shortcut menu for performing various functions on the selected signatures. Some commands appear only if you select a single signature, while some commands can be used on more than one signature at a time; your changes apply to all selected signatures. For more information about the Signatures policy, see Signatures Page, page 39-4.

Additionally, the available commands differ depending on which cell you right click. Some commands are available when right-clicking any cell, while others are specific to a single cell.

Tip When you use a right-click command to change the value in a cell of a Default signature, the signature is converted to a Local signature in Device view or a shared-policy-specific signature in Policy view.

The following table explains the available commands.

Table 39-3 Signature Shortcut Menu

Menu Command	Description
Commands Available for All Cells	
Add Row	Adds a custom signature. For more information, see the following topics: • Adding Custom Signatures, page 39-19 • Edit Signature or Add Custom Signature Dialog Boxes, page 39-15
Edit Row	Edits the selected signature. You can edit one signature at a time. For more information, see the following topics: • Editing Signatures, page 39-14 • Edit Signature or Add Custom Signature Dialog Boxes, page 39-15
Delete Row	Deletes the selected custom signatures. You cannot delete Cisco-defined signatures. If you do not want to deploy a Cisco-defined signature, you can retire it or disable it.
Clone	Creates a new custom signature that contains the same properties as the selected signature. For more information, see Cloning Signatures, page 39-21.
Enable, Disable	Places the signature in the enabled or disabled state, respectively. Disabled signatures appear with crosshatching over them. For more information, see Enabling and Disabling Signatures, page 39-14.
Show Events Show MARS Events	Enables navigation to the Event Viewer or Cisco Security MARS application to view the realtime or historical events detected by the selected signature. For more information, see Viewing Events for an IPS Signature, page 69-57 and Viewing CS-MARS Events for an IPS Signature, page 72-44.
Action Cell Commands	
Add to Actions	Adds an action to the current list of actions for the selected signature.
Delete from Actions	Deletes an action from the current list of actions for the selected signature.
Replace Actions With	Replace the current set of actions for the selected signature with the single action selected. If you want to select more than one action, select **More** from the submenu, then use Ctrl+click to select the desired actions.
Edit Actions	Opens the Edit Actions dialog box, where you can select the desired actions for the signature. Your selection replaces the current list of actions for the signature. For more information, see Edit, Add, Replace Action Dialog Boxes, page 39-12.

Table 39-3 Signature Shortcut Menu (continued)

Menu Command	Description
Severity Cell Commands	
• High • Medium • Low • Informational	Changes the severity level of the signature to the level you select.
Fidelity Cell Commands	
Edit Fidelity	Changes the fidelity rating of the signature, which is the weight associated with how well this signature might perform in the absence of specific knowledge of the target.
Retired Cell Commands	
• Retire • Activate • Retire on Low Memory • Retire on Medium Memory	Changes the retired status of the signature to the selected status. For more information about the retired status categories, see Edit Signature or Add Custom Signature Dialog Boxes, page 39-15.

Edit, Add, Replace Action Dialog Boxes

Use the Edit, Add, or Replace Action dialog boxes to change the actions defined for a signature. These dialog boxes are available only when you edit the Action cell using the right click menu as explained in Signature Shortcut Menu, page 39-10. The behavior differs depending on the dialog box name:

- Add Actions—The actions that you select are added to those already defined in the signature. To open this dialog box, right-click the Actions cell of a signature and select **Add to Actions > More**.
- Replace Actions—The actions that you select completely replace those defined in the signature. To open this dialog box, right-click the Actions cell of a signature and select **Replace Actions With > More**.
- Edit Actions—The actions that you select completely replace those defined in the signature. To open this dialog box, right-click the Actions cell of a signature and select **Edit Actions**.

For an explanation of the available actions, see Understanding IPS Event Actions, page 40-2. You can select multiple actions using Ctrl+click.

Note When you open dialog box, the list of actions that you see varies. The list of actions depends upon whether you right-click in only one signature row in the Actions column or select more than one signature row before right-clicking in the Actions column. If you right-click in only one signature row in the Actions column, the list of actions is that of the engine for that signature. If you select more than one signature row before right-clicking in the Actions column, the list of actions is that which is available for each affected engine. (It is the list of common actions, not the union of actions.)

Edit Fidelity Dialog Box

Use the Edit Fidelity dialog box to make changes in the Fidelity Rating for a particular signature. The Fidelity Rating, or Signature Fidelity Rating (SFR), identifies the weight associated with how well this signature might perform in the absence of specific knowledge of the target. This rating can be any number from 0 to 100, with 100 indicating the most confidence in the signature.

Navigation Path

In the Signatures policy, right-click the Fidelity cell in a signature and select Edit Fidelity. For information on opening the Signatures policy, see Signatures Page, page 39-4. For more information on the signature shortcut menu, see Signature Shortcut Menu, page 39-10.

Viewing Signature Update Levels

In Device view, you can determine the current signature update packages applied to the device in Security Manager and compare it to the one deployed on the device.

Differences between the applied and deployed update levels can occur when:

- The device is updated outside of Security Manager.
- An update is applied to the policy in Security Manager but not yet published to the device.
- During initial Security Manager deployment before the devices are under Security Manager control.

To view the signature update level, select the **IPS > Signatures > Signatures** policy for an IPS device in Device view. Then, click the **View Update Level** button to open the Update Level dialog box.

The following table describes the information displayed in the dialog box.

Table 39-4 Update Level Dialog Box

Element	Description
Applied Level	This column displays the patch level that is applied to this device in Security Manager.
Deployed Level	This column displays the patch level that is currently running on the selected device.
Major Update	Identifies the major update level.
Minor Update	Identifies the minor update level.
Service Pack	Identifies the service pack level.
Patch	Identifies the patch level.
Engine	Identifies the engine level.
Signature Update	Identifies the signature update level. **Note** This field is the only field on this page that applies to the IOS IPS devices; all of the other fields are exclusive to IPS devices.
Revert button	If you mistakenly modify Applied Level, allows you to discard that new Applied Level; clicking **Revert** syncs the Applied Level to the Deployed Level. **Tip** A warning dialog appears before performing Revert. Also, a warning dialog appears asking you to submit the activity.

Enabling and Disabling Signatures

You can enable and disable individual signatures. Your change takes effect when you redeploy the configuration to the device.

If a signature is disabled, it appears in the table overlain with hash marks. When you deploy the configuration, disabled signatures are removed from the device.

Disabling signatures is useful when you want to reduce the number of signatures used by a device, or if you want to temporarily stop using a custom signature without deleting it. You can later reenable a signature that you disabled.

Note You can enable a signature that is retired, but it then is not used to scan traffic, because it is not in the signature micro-engine. If you want a sensor to scan network traffic for a particular signature, you must enable it and not retire it. The AIP-SSC-5 does not support enabling a signature that is retired.

Step 1 Do one of the following:

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

The Signature page appears; see Signatures Page, page 39-4.

Step 2 Right-click the signature whose enabled status you want to change and select **Enable** or **Disable**, as appropriate.

Editing Signatures

You can edit signatures to change their behavior. For example, you can change the action that should be taken when a signature fires, or the severity and fidelity ratings used to calculate the risk rating of the signature.

Some signatures have special requirements. For example, to configure a sensor to detect ACL violation signatures, you must first configure one or more Cisco IOS routers to log ACL violations. Then, you must configure those routers to communicate with the sensor. Finally, you must configure the sensor to accept syslog traffic from those routers.

Tip This procedure describes how to edit an entire signature. You can also selectively edit individual properties of a signature using the right-click menu in the Signatures policy. For information on the available commands, see Signature Shortcut Menu, page 39-10.

Related Topics

- Understanding Signatures, page 39-1
- Understanding IPS Event Actions, page 40-2
- Enabling and Disabling Signatures, page 39-14

- Cloning Signatures, page 39-21
- Configuring Event Action Filters, page 40-4
- Chapter 40, "Configuring Event Action Rules"

Step 1 Do one of the following:

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

The Signature page appears; see Signatures Page, page 39-4.

Step 2 Right-click the signature you want to edit and select **Edit Row**. You can also select the signature and click the Edit Row (pencil) button beneath the signatures table. The Edit Signature dialog box opens.

Tip Use the filter fields above the table to help you find the desired signature. For information on filtering tables, see Filtering Tables, page 1-48.

Step 3 Make the desired changes to the signature. For specific details about each option, see Edit Signature or Add Custom Signature Dialog Boxes, page 39-15.

When editing signatures, keep the following in mind:

- You cannot edit a Default signature. Default signatures are the Cisco-defined version of a signature. Before you can edit a Default signature, you must convert it either to a Local signature (one defined specifically on the selected device) or a shared-policy-specific signature (one defined in a shared policy). You must select either Local or the shared policy name from the Source Policy field before you can change any field on the Edit Signature dialog box.
- You cannot change every characteristic of a signature. For example, you cannot change the signature or subsignature IDs. These fields are read-only.
- If you want to change the detailed parameters of a signature, follow the procedure described in Editing Signature Parameters (Tuning Signatures), page 39-23.

Step 4 Click **OK** to save your changes.

Edit Signature or Add Custom Signature Dialog Boxes

The Edit Signature and Add Custom Signature dialog boxes are essentially the same. Most of the fields are identical, although there are some layout differences. Use these dialog boxes as follows:

- Use the Edit Signature dialog box to edit the characteristics of a non-default signature (you can only view the characteristics of a default signature in read-only mode).

  You cannot edit default signatures. To make any changes to a signature, you must select something other than Default in the Source Policy field at the top of the dialog box.

- Use the Add Custom Signature dialog box to create a custom signature. In the Add Custom Signature dialog box, you enter a name and then select an existing engine from a drop-down list. The signature ID and subsignature ID are assigned by Security Manager. After you finish selecting the remaining parameters, the new signature is added to the Signatures page in the appropriate numerical location, and it is selected.

Note Beginning with Security Manager 4.4, you can specify a signature ID and a subsignature ID while adding a custom signature. If you specify a signatureID/subsignature ID combination that already exists, you will receive an error message.

Navigation Path

From the Signatures page:

- To edit a signature, right-click the policy you want to edit and select **Edit Row**.
- To add a custom signature, click the **Add Row (+)** button beneath the table, or right-click any row and select **Add Row**.

For information on opening the Signature page, see Signatures Page, page 39-4.

Related Topics

- Edit, Add, Replace Action Dialog Boxes, page 39-12
- Edit Signature Parameters Dialog Box, page 39-24
- Engine Options, page 39-20

Field Reference

Table 39-5 Edit Signature or Add Custom Signature Dialog Boxes

Element	Description
Source Policy (Edit signature only.)	The policy in which you are editing the signature: • Default—The default Cisco-defined signature, which you cannot edit. You must select something other than Default to edit the signature. • Local—The signature is a local signature defined specifically for the selected device. This option is not available in Policy view. • Policy name (variable)—The name of a shared policy. In Device view, a policy name is available only if you assign a shared policy to the device. In Policy view, this is the name of the policy you are editing. Select the policy name to edit the signature and to have your edits reflected on all devices that are assigned the shared policy.
Name (Add only.)	The name of the signature. You cannot change the name after you create the signature. If you want to change the name, you must create a clone of the signature.
SigID (Add only.)	The signature ID that you specify while adding a custom signature. The allowed range of values is 60000 - 65000.

Table 39-5 Edit Signature or Add Custom Signature Dialog Boxes (continued)

Element	Description
SubSigID (Add only.)	The subsignature ID that you specify while adding a custom signature. The allowed range of values is 0 - 255.
Inheritance Mandatory (Edit signature only.)	When selected, forces any policy that inherits from this policy to use the signature settings defined.
Enabled	Whether the signature is enabled.
Severity	The severity level that the signature will report: High, Medium, Low, or Informational.
Fidelity Rating	The weight associated with how well this signature might perform in the absence of specific knowledge of the target.
Actions	The actions that the sensor will take when this signature fires. For a complete list of actions, see the Understanding IPS Event Actions, page 40-2. Use Ctrl+click to select multiple actions.
Base Risk Rating Risk Rating (Fields have slightly different names when adding or editing signatures.)	The base risk rating value of the signature, which is calculated by multiplying the fidelity rating and the severity factor and dividing them by 100 (Fidelity Rating x Severity Factor /100). This value is read only; you cannot directly change it. To change the value, alter the Severity and Fidelity fields. The Severity Factor has the following values based on what you select in the Severity field: • High = 100 • Medium = 75 • Low = 50 • Informational = 25
Engine (Read-only when editing; read-write when adding custom signatures.)	The engine that parses and inspects the traffic specified by this signature. For a description of the engines, see Engine Options, page 39-20. When adding a custom signature, you must select the appropriate engine. For detailed information about each engine, and the parameters available, see the "Signature Engines" section in the *Installing and Using Cisco Intrusion Prevention System Device Manager* document for the IPS Software release you are using. **Tip** The term *engine* used here is not the same as the term *engine level* used in the row below.

Table 39-5 Edit Signature or Add Custom Signature Dialog Boxes (continued)

Element	Description
Retired	The conditions under which the signature is retired, if any. A retired signature is removed from the signature engine. You can activate a retired signature to place it back in the signature engine. **Timesaver** Use the retired field to unload disabled signatures on your IOS-IPS device to achieve the most favorable memory consumption of that device. If the engine level of a signature policy is less than E-4, the Retired field has two possible values: false and true. False means that the signature is not retired; true means that the signature is retired. If the engine level of a signature policy is equal to E-4, the Retired field has four possible values: • **false**—The signature is not retired. • **low-mem-retired**—The signature should be retired on low-memory platforms. A low-memory device is one that has 2 MB RAM or less. • **med-mem-retired**—The signature should be retired on both low-end and medium platforms. A medium-memory device is one that has 4 MB RAM or less, but more than 2 MB RAM. (Any device with more than 4 MB RAM is considered a high-memory platform.) • **true**—The signature is retired on all platforms. When you select low-mem-retired or med-mem-retired, Security Manager configures the device with those signatures. Whether the signature is actually retired on the device depends on amount of memory installed on the device; the device makes the decision on which signatures are actually retired. **Tip** The term *engine level* used here is not the same as the term *engine* in the row above.
Obsolete (Edit signature only.)	Identifies whether the signature is obsolete. An obsolete signature is removed from the signature engine. It cannot be re-activated.
Restore Defaults button (Non-custom signatures only. Edit signature only.)	Click this button to revert to default values for this signature as defined by Cisco.
Edit Parameters button	Click this button to edit the detailed parameters for this signature using the Edit Signature Parameters dialog box. For more information, see the following topics: • Edit Signature Parameters Dialog Box, page 39-24 • Editing Signature Parameters (Tuning Signatures), page 39-23

Adding Custom Signatures

If you want to look for traffic patterns that are not identified by the built-in signatures, you can create your own custom signatures to define the traffic patterns.

Even if a built-in signature covers the traffic pattern, you might want to create a custom signature to edit the detailed signature parameters without altering the default signature. If you are creating a custom signature that is similar to an existing signature, the easiest way to do it is to clone the signature as described in Cloning Signatures, page 39-21.

When adding a custom signature to some IPS devices, you can use regular expressions. For more information on the importance of proper syntax when using regular expressions, refer to Regular Expressions in Custom Signatures, page 39-22

Note The AIP-SSC-5 does not support custom signatures.

Step 1 Do one of the following:

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

The Signature page appears; see Signatures Page, page 39-4.

Step 2 Click the **Add Row (+)** button beneath the signature table to open the Add Custom Signature dialog box.

Step 3 Configure the desired settings. For specific details about each option, see Edit Signature or Add Custom Signature Dialog Boxes, page 39-15.

When creating signatures, keep the following in mind:

- You cannot change the signature name after you define the signature. If you later want to change the name, you must clone the signature and change it while creating the clone.
- Select the appropriate signature engine for the signature. For a description of the signature engines, see Engine Options, page 39-20. You cannot change the engine after you create the signature; if you select the wrong engine and click OK to save the signature, you must start over and create an entirely new signature.
- The default is to create an enabled signature, but you can deselect the Enabled check box to initially create a disabled signature. You might want to disable the signature if you have not finished editing its parameters.
- Follow the procedure described in Editing Signature Parameters (Tuning Signatures), page 39-23 to define the detailed signature parameters. You must select the desired engine before you edit the parameters, because many of the parameters are determined by the signature engine.

  Whether you can save the signature before configuring parameters differs based on the engine that you select. At minimum, you must click **Edit Parameters** to open the Edit Signature Parameters dialog box, and then click **OK** in the Edit Signature Parameters dialog box, before you can save the signature definition. However, to create a meaningful signature, you will need to configure the parameters to identify the desired traffic pattern.

Step 4 Click **OK** to save your changes.

The custom signature is added to the end of the table and given the next available signature ID starting at 60000.

Note Beginning with Security Manager 4.4, you can specify a signature ID and a subsignature ID while adding a custom signature. If you specify a signatureID/subsignature ID combination that already exists, you will receive an error message.

Engine Options

The following list identifies the options you can specifying in the Engine field of the Edit Signature Parameters dialog box. For detailed information about each engine, and the parameters available, see the "Signature Engines" section in the *Installing and Using Cisco Intrusion Prevention System Device Manager* document for the IPS Software release you are using.

- AIC FTP—Inspects FTP traffic and lets you control the commands being issued.
- AIC HTTP—Provides granular control over HTTP sessions to prevent abuse of the HTTP protocol.
- Atomic ARP—Inspects Layer-2 ARP protocol. The Atomic ARP engine is different because most engines are based on Layer-3-IP.
- atomic-ip—Inspects IP protocol packets and associated Layer-4 transport protocols.
- Atomic IPv6—Detects IOS vulnerabilities that are stimulated by malformed IPv6 traffic.
- Flood Host—Detects ICMP and UDP floods directed at hosts.
- Flood Net—Detects ICMP and UDP floods directed at networks.
- Meta—Defines events that occur in a related manner within a sliding time interval. This engine processes events rather than packets.
- multi-string—Defines signatures that inspect Layer 4 transport protocol (ICMP, TCP, and UDP) payloads using multiple string matches for one signature. You can specify a series of regular expression patterns that must be matched to fire the signature.
- normalizer—Configures how the IP and TCP normalizer functions and provides configuration for signature events related to the IP and TCP normalizer. Allows you to enforce RFC compliance.
- service-dns—Inspects DNS (TCP and UDP) traffic.
- service-ftp—Inspects FTP traffic.
- Service Generic—Decodes custom service and payload.

  The Service Generic engine allows programmatic signatures to be issued in a config-file-only signature update. It has a simple machine and assembly language that is defined in the configuration file. It runs the machine code (distilled from the assembly language) through its virtual machine, which processes the instructions and pulls the important pieces of information out of the packet and runs them through the comparisons and operations specified in the machine code. It is intended as a rapid signature response engine to supplement the String and State engines.

  You cannot use the Service Generic engine to create custom signatures.

  Note Due to the proprietary nature of this complex language, we do not recommend that you edit the Service Generic engine signature parameters. Change only the severity and event action for these signatures.

- Service Generic Advanced—Generically analyzes network protocols.
- Service H225—Inspects VoIP traffic.
- service-http—Inspects HTTP traffic. The WEBPORTS variable defines inspection port for HTTP traffic.
- Service IDENT—Inspects IDENT (client and server) traffic.
- Service MSRPC—Inspects MSRPC traffic.
- Service MSSQL—Inspects Microsoft SQL traffic.
- Service NTP—Inspects NTP traffic.
- service-rpc—Inspects RPC traffic.
- Service SMB—Inspects SMB traffic.
- Service SMB Advanced—Processes Microsoft SMB and Microsoft RPC over SMB packets.
- Service SNMP—Inspects SNMP traffic.
- Service SSH—Inspects SSH traffic.
- Service TNS—Inspects TNS traffic.
- state—Stateful searches of strings in protocols such as SMTP.
- string-icmp—Searches on Regex strings based on ICMP protocol.
- string-tcp—Searches on Regex strings based on TCP protocol.
- string-udp—Searches on Regex strings based on UDP protocol.
- Sweep—Analyzes sweeps of ports, hosts, and services, from a single host (ICMP and TCP), from destination ports (TCP and UDP), and multiple ports with RPC requests between two nodes.
- Sweep Other TCP—Analyzes TCP flag combinations from reconnaissance scans that are trying to get information about a single host. The signatures look for flags A, B, and C. When all three are seen, an alert is fired.
- Traffic ICMP—Analyzes nonstandard protocols, such as TFN2K, LOKI, and DDOS. There are only two signatures with configurable parameters.
- Traffic Anomaly—Analyzes TCP, UDP, and other traffic for worm-infested hosts.
- Trojan Bo2k—Analyzes traffic from the nonstandard protocol BO2K. There are no user-configurable parameters in this engine.
- Trojan Tfn2k—Analyzes traffic from the nonstandard protocol TFN2K. There are no user-configurable parameters in this engine.
- Trojan UDP—Analyzes traffic from the UDP protocol. There are no user-configurable parameters in this engine.

Cloning Signatures

If you want to create a custom signature that is similar to an existing signature, you can create a clone, or copy, of the signature. You can then edit the parameters to make the clone perform according to your requirements.

For example, you might want to create a clone of a Cisco-defined signature to customize it to your needs. You might find this preferable to converting the Cisco signature to a Local or shared policy signature and directly editing its parameters.

To clone a signature, follow these steps:

Step 1 Do one of the following:

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

The Signature page appears; see Signatures Page, page 39-4.

Step 2 Right-click the signature that you want to clone and select **Clone**.

Security Manager takes some time to make the copy, and might warn you that some attributes are read-only and cannot be copied. If you receive a warning, click **OK**. The Add Custom Signature dialog box then appears.

Step 3 Edit the properties of the cloned signature, as described in Adding Custom Signatures, page 39-19.

Step 4 Click **OK**. The clone appears in the summary table on the Signatures page as the last signature.

Cloned signatures are enabled and active by default.

Regular Expressions in Custom Signatures

You can use regular expressions when adding a custom signature to some IPS devices.

Regardless of the type of IPS device or the particular characteristics of the custom signature, incorrect syntax in a regular expression will cause device deployment to fail after adding the custom signature.

Regular expressions may contain many control characters or regex notations which are used to describe the regex pattern itself. If you intend to use them as literal characters in the regular expression itself, they should be escaped with the "\" escape character. If you intend to use them for their real meaning, on the other hand, you should be careful to conform to proper regular expression syntax.

Example of regular expression that will cause deployment failure: !@#$%^&*()_+{}|:"<>?

Example of regular expression that will be deployed successfully: !@#$%\^&*\(\)_+\{\}|:\"<>\?

Using regular expressions in custom IPS signatures is described in this example:

Step 1 Add a Cisco ASA 5500 Series IPS Security Services Processor (e.g., 5525-X).

Step 2 Add a custom signature with a string-XL engine (e.g., string-xl-tcp) to the IPS device.

Step 3 Click Edit Parameter and create a regular expression for the custom signature.

Step 4 Deploy the IPS device.

Step 5 Deployment will fail if incorrect syntax for regular expressions used, but deployment will succeed if correct syntax is used.

Editing Signature Parameters (Tuning Signatures)

If you cannot alter the behavior of a signature to fit your needs using the Event Action Filters and Overrides policies, or by changing the actions associated with a signature, you might need to fine-tune the signature parameters. You should considered editing parameters to be your last option, however, because these parameters can be complex and frequently require that you have a deep understanding of packet characteristics.

The reason you would want to edit parameters is to reduce false positives and false negatives:

- A *false positive* occurs when legitimate network activity, such as virus scanning, is interpreted and reported as an attack. This happens when network activity meets criteria that were specified to identify an attack before the attack occurred. You can decrease the number false positives by tuning your sensor configurations.
- A *false negative* occurs when an attack was not detected. Tuning your sensor configurations will help you decrease the number of false negatives.

Tip You cannot edit the parameters of a default signature. Before editing the parameters of a default signature, you must convert the signature to a local- or shared-policy signature. In some cases, such as regular expression editing, you must clone the signature and convert it to a custom signature.

This procedure describes how to edit signature parameters to tune a signature.

Related Topics

- Editing Signatures, page 39-14
- Understanding Signatures, page 39-1
- Configuring Event Action Filters, page 40-4
- Configuring Event Action Overrides, page 40-13

Step 1 Do one of the following:

- (Device view) Select **IPS > Signatures > Signatures** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Signatures > Signatures**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Signatures**, then select an existing policy or create a new one.

The Signature page appears; see Signatures Page, page 39-4.

Step 2 Right-click the signature whose parameters you want to edit and select **Edit Row**. The Edit Signature dialog box appears (see Edit Signature or Add Custom Signature Dialog Boxes, page 39-15).

Step 3 If the Source Policy field shows Default, you must change it to Local or to the name of a shared policy before you can edit the parameters. The Local option is available in Device view only, and makes your changes apply to the device you are editing and to no other devices. If you select the name of a shared policy, your changes apply to all devices that are assigned the policy.

Step 4 Click **Edit Parameters**. The Edit Signature Parameters dialog box appears.

The Edit Signature Parameters dialog box contains a folder tree structure, with the parameter names in the left side tree, and the values of the parameters shown on the right side.

Values that you can change contain a little box in the name; this is a check box. An empty check box indicates that the default value is being used for the parameter. Check the check box to configure that parameter. Click the value field to change the parameter. A green check indicates that a user-defined value is being used. Click the green check to change the value back to the default. (Editing the field typically adds a check mark to the box.)

To change a parameter, click in the associated filed in the right side. The behavior of clicking on a parameter differs based on the parameter type:

- Read-only parameters—Many parameters are read-only and cannot be changed, such as signature ID. Clicking these parameters typically has no effect, although parameter lists will open a dialog box (such as the Obsoletes list).
- Text or Numeric parameters—When you click a parameter that requires that you type in a value, whether alphanumeric or numeric, the field becomes an edit box. Type in the desired value and either press enter or click outside the edit box.
- Predefined value parameters—Many parameters have a small set of possible values, such as Yes/No. When you click these parameters, you activate a drop-down list. Select the desired option and click outside the field.
- List parameters—Some parameters contain a list of items. These parameters are represented by a pencil icon in the parameter value along with a word, such as Set or List. When you click in the field, a dialog box opens where you can configure the list associated with the item. The Meta engine component list is an example; for more information, see Editing the Component List for Meta Engine Signatures, page 39-29.
- Variable parameters—Some parameters allow you to select policy objects to identify the contents of the parameters. For example, you can select port list objects to identify ports in some signature engines. When you click these parameters, an edit box with a Select button appears. You can type the items directly into the edit box, including the name of the policy object, or click **Select** to select the policy object from a list or to create a new object.

For more information about the Edit Signature Parameters dialog box, see Edit Signature Parameters Dialog Box, page 39-24.

Step 5 Change the settings as desired, then click **OK** to save your changes. You are returned to the Edit Signature dialog box.

Step 6 Click **OK** in the Edit Signature dialog box to save your changes to the signature.

Tip If you decide that your edits did not have the desired effect, or you suspect that you made a mistake, you can click the **Restore Defaults** button in the Edit Signature dialog box to erase your changes. You can then start over.

Edit Signature Parameters Dialog Box

Use the Edit Signature Parameters dialog box to edit (also called tune) the built-in micro-engine parameters for a particular signature. Different engines have different parameters, so the appearance of the Edit Signature Parameters dialog box varies. For more information about editing signature parameters, see Editing Signature Parameters (Tuning Signatures), page 39-23.

The Edit Signature Parameters dialog box contains a folder tree structure, with the parameter names in the left side tree, and the values of the parameters shown on the right side.

Values that you can change contain a little box in the name; this is a check box. An empty check box indicates that the default value is being used for the parameter. Check the check box to configure that parameter. Click the value field to change the parameter. A green check indicates that a user-defined value is being used. Click the green check to change the value back to the default. (Editing the field typically adds a check mark to the box.)

To change a parameter, click in the associated filed in the right side. The behavior of clicking on a parameter differs based on the parameter type:

- Read-only parameters—Many parameters are read-only and cannot be changed, such as signature ID. Clicking these parameters typically has no effect, although parameter lists will open a dialog box (such as the Obsoletes list).
- Text or Numeric parameters—When you click a parameter that requires that you type in a value, whether alphanumeric or numeric, the field becomes an edit box. Type in the desired value and either press enter or click outside the edit box.
- Predefined value parameters—Many parameters have a small set of possible values, such as Yes/No. When you click these parameters, you activate a drop-down list. Select the desired option and click outside the field.
- List parameters—Some parameters contain a list of items. These parameters are represented by a pencil icon in the parameter value along with a word, such as Set or List. When you click in the field, a dialog box opens where you can configure the list associated with the item. The Meta engine component list is an example; for more information, see Editing the Component List for Meta Engine Signatures, page 39-29.
- Variable parameters—Some parameters allow you to select policy objects to identify the contents of the parameters. For example, you can select port list objects to identify ports in some signature engines. When you click these parameters, an edit box with a Select button appears. You can type the items directly into the edit box, including the name of the policy object, or click **Select** to select the policy object from a list or to create a new object.

Navigation Path

From the Edit Signature or Add Custom Signature dialog boxes, click the **Edit Parameters** button. For information on opening these dialog boxes, see Edit Signature or Add Custom Signature Dialog Boxes, page 39-15.

Tip If the button is not active, you must first select Local or the name of a shared policy from the Source Policy field, or clone the signature to create a custom policy. The Local option is available in Device view only, and makes your changes apply to the device you are editing and to no other devices. If you select the name of a shared policy, your changes apply to all devices that are assigned the policy.

Field Reference

Table 39-6 Edit Signature Parameters Dialog Box

Elements	Description
Tuning Context (Policy view only)	Displays the information needed by Security Manager to uniquely describe how the signature parameters were edited (tuned) for a particular signature policy. The Tuning Context field is a character string that contains the following items: • **Context**—The identification given by the Security Manager server for the unique definition of this micro-engine. • **SigLevel (IPS) or Version (IOS IPS)**—For IPS policies, the range of signature update levels to which this definition of the signature micro-engine applies. For IOS IPS, this is the IOS IPS version. • **Engine**—The name of the IPS engine. **Tip** As an example, the Tuning Context field could contain the following character string: Context:9, SigLevel:302-449, Engine:atomic-ip. For any particular signature policy, the Tuning Context field can contain one or many tuning contexts: • The tuning context with the highest signature level is pre-pended with "Reference context." • If you modify the shared policy that is pre-pended with "Reference context," Security Manager may ask you if you want to copy the policy to other applicable contexts. (A particular device may appear in more than one context.) • If you choose to copy the policy to other applicable contexts, an error message informs you if some parameters cannot be copied. **Note** Beginning with Security Manager 4.1, old signature versions (defined as being older than the lowest signature level deployed) are purged during a periodic purge operation, the purpose of which is to optimize the database. As a result, you may notice the deletion of some of your unused tuning contexts.
Signature ID	The unique numerical value assigned to this signature. This value lets the sensor identify a particular signature. The value is 1000 to 65000.
SubSignature ID	The unique numerical value assigned to this subsignature. The subsignature ID identifies a more granular version of a broad signature. The value is 0 to 255.
Promiscuous Delta	Modifies the seriousness of an alert when operating in promiscuous mode. The value is subtracted from an alert's overall risk rating. The promiscuous delta is ignored when operating in inline mode. The value can be 0 to 30.

Table 39-6 Edit Signature Parameters Dialog Box (continued)

Elements	Description
Sig Description	A description of the signature to help you distinguish this signature from other signatures: • Alert Notes—Additional information about this signature that will be included in the alert message. • User Comments—Your comments about the signature. • Alarm Traits—Traits you want to document about this signature. The value is 0 to 65535. The default is 0. • Release—The release in which the signature was most recently updated. • Signature Creation Date—The date on which the signature was created. • Signature Type—The type of signature: Anomaly, Component, Exploit, Vulnerability, or Other.
Engine	The engine that parses and inspects the traffic specified by this signature. The engine determines which parameters are available in the Engines folder. For a description of the engines, see Engine Options, page 39-20. For detailed information about each engine, and the parameters available, see the "Signature Engines" section in the *Installing and Using Cisco Intrusion Prevention System Device Manager* document for the IPS Software release you are using. **Tip** Many engines include the Fragment Status parameter, which lets you identify whether packet fragments should be inspected. You can elect to not inspect fragments, to inspect fragments, or to apply the signature to any packet status.
Event Counter	How the sensor counts events. For example, you can specify that you want the sensor to send an alert only if the same signature fires 5 times for the same address set. Configure the following values: • Event Count—The number of times an event must occur before an alert is generated. The value is 1 to 65535. The default is 1. • Event Count Key—The storage type used to count events for this signature. Choose attacker address, attacker address and victim port, attacker and victim addresses, attacker and victim addresses and ports, or victim address. The default is attacker address. • Specify Alert Interval—Whether you want to specify the time between alerts for resetting the event count, Yes or No. If you select Yes, enter the time in seconds from 2 to 1000.

Table 39-6 Edit Signature Parameters Dialog Box (continued)

Elements	Description
Alert Frequency	How often the sensor alerts you when this signature is firing. Specify the following parameters for this signature. These parameters are explained below. • Summary Mode • Summary Interval • Summary Key • Specify Global Summary Threshold
Summary Mode (Alert Frequency group)	The mode of alert summarization. There are four modes: Fire All, Fire Once, Summarize, and Global Summarize. The summary mode is changed dynamically to adapt to the current alert volume. For example, you can configure the signature to Fire All, but after a certain threshold is reached, it starts summarizing. Your selection of summary mode controls which other parameters are available in the Summary Mode group. • Fire All—Fires an alert on all events. • Fire Once—Fires an alert only once. • Summarize—Summarizes alerts. • Global Summarize—Summarizes an alert so that it only fires once regardless of how many attackers or victims. **Note** When multiple contexts from an ASA device are contained in one virtual sensor, the summary alerts contain the context name of the last context that was summarized. Thus, the summary is the result of all alerts of this type from all contexts that are being summarized.
Specify Summary Threshold (Summary Mode group.)	When you select Fire All, you can select whether you want to configure the summary threshold settings that will be used if the device dynamically changes to summary mode. If you select Yes, you can configure the summary interval, key, or global summary thresholds.
Summary Interval (Summary Mode group.)	The time in seconds used in each summary alert. The value is 1 to 65535. The default is 15.
Summary Key (Summary Mode group.)	The storage type used to summarize alerts. Choose Attacker address, Attacker address and victim port, Attacker and victim addresses, Attacker and victim addresses and ports, or Victim address. The default is Attacker address.
Specify Global Summary Threshold (Summary Mode group.)	Whether to specify the threshold number of events to take the alert into global summary, Yes or No. If you select Yes, enter the threshold number of events, from 1 to 65535. The default is 240.

Table 39-6 Edit Signature Parameters Dialog Box (continued)

Elements	Description
Status	The status of the signature. The Obsoletes list shows the signatures that are obsoleted by this signature; click the pencil icon to open the list. In many cases, this information is read-only. If you can modify the list, click Set in the parameter field to open the list, where you can add the obsoleted signature IDs.
Vulnerable OS List	The list of operating systems that this attack targets.
MARS Category	The category in Cisco Security MARS to which this signature belongs. This metadata is used to color the events generated in such a way as to provide MARS with the data that it needs to process this signature relative to the event categories that it studies.
Expand All button	Expands all categories and subcategories.
Collapse All button	Collapses all fields to the category.

Editing the Component List for Meta Engine Signatures

Use the Edit Signature Parameter—Component List dialog box to edit the component list for a meta engine signature.

The Meta engine defines events that occur in a related manner within a sliding time interval. This engine processes events rather than packets. As signature events are generated, the Meta engine inspects them to determine if they match any or several Meta definitions. The Meta engine generates a signature event after all requirements for the event are met.

All signature events are handed off to the Meta engine by the Signature Event Action Processor. The Signature Event Action Processor hands off the event after processing the minimum hits option. Summarization and event action are processed after the Meta engine has processed the component events.

The Meta engine is different from other engines in that it takes alerts as input where most engines take packets as input. Thus, in a Meta engine signature, you must identify the signatures that the Meta signature should be looking for. This list of signatures is contained in the Component list.

The Component list is part of the signature parameters. To edit the parameters, follow the procedure described in Editing Signature Parameters (Tuning Signatures), page 39-23. When you open the Edit Signature Parameters dialog box for a signature that uses the Meta engine, look for the **Engine > Component List** parameter. The parameter value contains a pencil icon and the word List. Click **List** to open the Edit Signature Parameter—Component List dialog box.

The dialog box is divided into two lists, an Inactive list (on the left) and an active list (on the right). The active list defines the signatures that the Meta engine signature is looking for.

To modify the components list:

- **Add new components**—Click the **Add Entry (+)** button to the left of the inactive list. The Add Signature Parameter—List Entry dialog box opens. Configure the following values:
  - **Entry Key**—A name for the component.
  - **Component Sig ID**—The signature ID of the signature you are looking for.
  - **Component SubSig ID**—The subsignature ID; enter 0 if there are no subsignatures.

  - **Component Count**—The number of times this signature must fire before the Meta signature is triggered.
  - **Is a Not Component**—This field lets you create negative entries; thus, you can identify a list where some signatures must fire, and some signatures must not fire. Select **No** for signatures that must fire, and **Yes** for signatures that must not fire.

  When you click **OK** in the Add Signature Parameter—List Entry dialog box, the new component is added to the inactive list. Select it and click the >> button to move it to the active list. Then, use the Up and Down arrow buttons to position the component in the active component list; a third button is available to reset the order to the previously saved order.

- **Edit an existing component**—Select the component (in either list) and click the **Edit Entry (pencil)** button that is between the lists. The Edit Signature Parameter—List Entry dialog box opens. The parameters are the same as for adding a new entry, except that you cannot change the component name.
- **Delete a component**—Select the component in the inactive list and click the **Delete Entry (trash can)** button that is to the left of the inactive list. If you want to delete an active component, you must first select it in the active list and click the << button to move it to the inactive list.
- **Restore defaults**—If you want to restore the default values of a component, select it and click **Restore**.

Obsoletes Dialog Box

Use the Obsoletes dialog box to identify obsolete signatures associated with a particular signature. In many cases, this information is read-only. In some cases, it is read-write; for example, you can edit the list for IOS IPS signature policies for Local or shared-policy-specific signatures.

If you can edit the list:

- Click the **Add Entry (+)** button to add the signature and subsignature ID of a signature that is made obsolete by the signature you are editing.
- Select an entry and click the **Delete Entry (trash can)** button to remove it from the list of obsoleted signatures.

Navigation Path

The Obsoletes list is part of the signature parameters. To edit the parameters, follow the procedure described in Editing Signature Parameters (Tuning Signatures), page 39-23. When you open the Edit Signature Parameters dialog box, look for the **Status > Obsoletes** parameter. The parameter value contains a pencil icon and the word Set (when the parameter is not read-only). Click the pencil or word to open the Obsoletes dialog box.

Configuring Signature Settings

Use the Signature Settings page to define settings for IPS appliances and service modules (but not Cisco IOS IPS devices). These settings define the following policies:

- **Application policy**—Enable or disable HTTP, determine and specify the maximum number of HTTP requests, specify AIC web ports, and enable or disable FTP.
- **Fragment reassembly policy**—Configure the sensor to reassemble a datagram that has been fragmented over more than one packet by selecting the IP reassembly mode.

- **Stream reassembly policy**—Configure the sensor to monitor only TCP sessions that have been established by a complete three-way handshake by specifying whether a TCP handshake is required and by selecting the TCP reassembly mode.
- **IP logging policy**—Configure the sensor to generate an IP session log when the sensor detects an attack by determining and selecting the maximum allowable number of log packets, the IP log time and the maximum allowable size of the IP log.

Tip All of these settings have default values, so configure this policy only if you need to use a non-default value.

To configure the Signature Settings policy, do one of the following:

- (Device view) Select **IPS > Signatures > Settings** from the Policy selector.
- (Policy view) Select **IPS > Signatures > Settings**, then select an existing policy or create a new one.

You can then configure the options that are explained in the following table.

Table 39-7 Signature Settings Page

Element	Description
Enable HTTP	Enables protection for web services. Select **Yes** to require the sensor to inspect HTTP traffic for compliance with the RFC.
Max HTTP Requests	The maximum number of outstanding HTTP requests per connection.
AIC Web Ports	The ports on which to look for AIC traffic. Enter a comma-separated list of port numbers or port list objects that define the ports. You can click **Select** to select a port list object from a list or to create a new object.
Enable FTP	Enables protection for FTP services. Select **Yes** to require the sensor to inspect FTP traffic.
IP Reassembly Mode	The method the sensor uses to reassemble the fragments, based on the operating system.
TCP Handshake Required	Whether the sensor should only track sessions for which the three-way handshake is completed.
TCP Reassembly Mode	The mode the sensor should use to reassemble TCP sessions with the following options: • Asymmetric—May only be seeing one direction of bidirectional traffic flow. **Note** Asymmetric mode lets the sensor synchronize state with the flow and maintain inspection for those engines that do not require both directions. Asymmetric mode lowers security because full protection requires both sides of traffic to be seen. • Loose—Use in environments where packets might be dropped. • Strict—If a packet is missed for any reason, all packets after the missed packet are not processed.
Max IP Log Packets	The number of packets you want logged.
IP Log Time	The duration you want the sensor to log, from 1 to 60 minutes. The default is 30 minutes.

Table 39-7 Signature Settings Page (continued)

Element	Description
Max IP Log Bytes	The maximum number of bytes you want logged.

Configuring Event Action Rules

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

An IPS event is an IPS message that contains an alert, a block request, a status message, or an error message. An event action is the sensor's response to an event. An event action happens only if the event is not filtered. Possible event actions are TCP reset, block host, block connection, IP logging, and capturing the alert trigger packet. Event actions were known as alarms in Cisco IPS versions earlier than 5.x.

The IPS Event Actions folder is where you configure settings for the event action processing component of the sensor. These settings define the actions for the sensor to take when an event is detected.

Note You cannot use IPv6 addresses in Event Action policies in Security Manager. For more information on IPv6 support in Security Manager, see IPv6 Support in Security Manager, page 1-8.

This chapter contains the following topics:

- Understanding the IPS Event Action Process, page 40-1
- Understanding IPS Event Actions, page 40-2
- Configuring Event Action Filters, page 40-4
- Configuring Event Action Overrides, page 40-13
- Configuring IPS Event Action Network Information, page 40-17
- Configuring Settings for Event Actions, page 40-23

Understanding the IPS Event Action Process

The IPS event action rules dictate the actions that the sensor performs when an event occurs. Although each signature is configured with specific actions that should be taken, the actual actions performed also depend on other factors.

Following is the general process that occurs when inspection identifies a signature event:

1. A signature alert occurs with actions specified by the signature. A risk rating for the alert is calculated.

For a detailed explanation of how risk rating is calculated, see Calculating the Risk Rating in *Installing and Using Cisco Intrusion Prevention System Device Manager 7.0* on Cisco.com.

You can influence risk ratings by configuring target value ratings and OS mappings; see Configuring IPS Event Action Network Information, page 40-17.

2. The **Event Action Overrides** policy is processed. If the risk rating of the event matches an override rule, the actions identified in the override rule are **added** to the actions defined in the signature. The overrides do not replace the actions specified in the signature.

   For information on configuring overrides, see Configuring Event Action Overrides, page 40-13.

3. The **Event Action Filters** policy is processed. If rules apply to the event, the rules **subtract** actions from the event. Thus, an action you added in a signature policy or override rule might be removed by one of your filter rules.

   For information on creating filter rules, see Configuring Event Action Filters, page 40-4.

4. Event summarization occurs, unless you turn off the summarization feature as described in Configuring Settings for Event Actions, page 40-23.

5. The actions are performed. For an explanation of possible actions, see Edit, Add, Replace Action Dialog Boxes, page 39-12.

6. A list of denied attackers is maintained, and subsequent access prevented, based on configurable settings. To change the default settings, see Configuring Settings for Event Actions, page 40-23.

Understanding IPS Event Actions

When you configure an event action filter or override, or a signature, you specify an action for events that meet the rule. For signatures and overrides, you are specifying an action to add to the event; for filters, you are specifying an action to remove from the event.

The most common action is Produce Alert, which generates an alert that you can view in your network management system, such as the Security Manager Event Viewer or CS MARS. However, there are a wide variety of actions that you can assign to an event. When looking over the possible actions, keep the following in mind:

- Many actions produce alerts in addition to the other action performed. The description for each action explains whether an alert is also produced.
- Cisco IOS IPS supports fewer actions for event action override or filter rules. The actions supported are Deny Attacker Inline, Deny Connection Inline, Deny Packet Inline, Product Alert, and Reset TCP Connection.
- Not all actions are necessarily available on all combinations of IPS software version and device type. Whenever you need to select an action, only those actions that are valid are available for selection.
- For deny and block actions, use the event actions settings policy to set the period of time for which addresses or packets are denied. For more information, see Configuring Settings for Event Actions, page 40-23.

The following table explains the possible actions.

Table 40-1 IPS Event Actions

Menu Command	Description
Deny Attacker Inline	Terminates the current packet and future packets from this attacker address for a specified period of time. The IPS must be operating in inline mode. For Cisco IOS IPS devices, no connection can be established from the attacker to the router until the shun time expires. **Tip** This is the most severe of the deny actions. It denies current and future packets from a single attacker address. For IPS appliances and service modules, you can use the IPS Device Manager to see a list of denied attackers and clear the list if necessary.
Deny Attacker/Service Pair Inline	Does not transmit this packet and future packets on the attacker address victim port pair for a specified period of time. The IPS must be operating in inline mode.
Deny Attacker/Victim Pair Inline	Does not transmit this packet and future packets on the attacker/victim address pair for a specified period of time. The IPS must be operating in inline mode.
Deny Connection Inline	Terminates the current packet and future packets on this TCP flow. Other connections from the attacker can be established. The IPS must be operating in inline mode.
Deny Packet Inline	Terminates the packet. The IPS must be operating in inline mode. For Cisco IOS IPS devices, this action discards the packet without sending a reset. Cisco recommends using “drop and reset” in conjunction with alarm. **Tip** For IPS appliances and service modules, there is an event action override that adds this action to high risk events. You cannot delete the override. If you do not want to use it, disable the override. For more information, see Configuring Event Action Overrides, page 40-13.
Log Attacker Packets	Starts IP logging on packets that contain the attacker address and sends an alert. This action causes an alert to be written to the Event Store, even if Produce Alert is not selected.
Log Pair Packets	Starts IP Logging on packets that contain the attacker/victim address pair. This action causes an alert to be written to the Event Store, even if Produce Alert is not selected.
Log Victim Packets	Starts IP Logging on packets that contain the victim address and sends an alert. This action causes an alert to be written to the Event Store, even if Produce Alert is not selected.
Modify Packet Inline	Modifies packet data to remove ambiguity about what the endpoint might do with the packet. **Tip** This option is not available for event action override or filter rules. It is available in signatures.

Table 40-1 IPS Event Actions (continued)

Menu Command	Description
Product Alert	Writes the event to the Event Store as an alert. For Cisco IOS IPS devices, the notification is sent through syslog or SDEE. **Note** A Produce Alert event action is added for an event when global correlation has increased the risk rating of an event, and has added either the Deny Packet Inline or Deny Attacker Inline event action.
Produce Verbose Alert	Includes an encoded dump of the offending packet in the alert. This action causes an alert to be written to the Event Store, even if Produce Alert is not selected.
Request Block Connection	Sends a request to block this connection. You must have blocking devices configured to implement this action. For more information, see Configuring IPS Blocking and Rate Limiting, page 43-7.
Request Block Host	Sends a request to block this attacker host. You must have blocking devices configured to implement this action.
Request Rate Limit	Sends a rate limit request to perform rate limiting. You must have rate limiting devices configured to implement this action.
Request SNMP Trap	Requests that the sensor send an SNMP trap notification to the configured trap destinations. This action causes an alert to be written even if Produce Alert is not selected. You must have SNMP configured on the sensor for traps to actually be sent. For more information, see Configuring SNMP, page 36-8.
Reset TCP Connection	Sends TCP resets to hijack and terminate the TCP flow, sending a reset to both the source and destination addresses. Reset TCP Connection works only on TCP signatures that analyze a single connection, for example, half-open SYN attacks. It does not work for sweeps or floods.

Related Topics

- Configuring Event Action Filters, page 40-4
- Configuring Event Action Overrides, page 40-13
- Configuring Signatures, page 39-4

Configuring Event Action Filters

You can configure event action filters to remove specific actions from an event or to discard an entire event and prevent further processing by the sensor.

Filters let the sensor perform certain actions in response to the event without requiring the sensor to perform all actions or remove the entire event. Filters work by removing actions from an event. A filter that removes all actions from an event effectively consumes the event. Before configuring filter rules, read Tips for Managing Event Action Filter Rules, page 40-6.

Note When filtering sweep signatures, we recommend that you do not filter the destination addresses. If there are multiple destination addresses, only the last address is used for matching the filter.

Related Topics

- Understanding the IPS Event Action Process, page 40-1

Step 1 Do one of the following to open the Event Action Filters policy:

- (Device view) Select **IPS > Event Actions > Event Action Filters** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Event Action Filters**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Event Action Filters**, then select an existing policy or create a new one.

The table shows the existing filter rules organized into sections. The Local section is for rules defined specifically for a selected device (in Device view). For shared or inherited policies, there are also sections for mandatory and default rules. For more information about the contents of this policy, see Event Action Filters Page, page 40-7.

Step 2 Select the row after which you want to create the filter rule and click the **Add Row** button or right-click and select **Add Row**. This opens the Add Filter Item dialog box. For detailed information about the options in this dialog box, see Filter Item Dialog Box, page 40-9.

Tips

- If you do not select a row, the new rule is added at the end of the local scope.
- You can also select an existing row and edit either the entire row (by clicking the **Edit Row** button) or specific cells. To edit a specific cell, right-click the cell and select the **Edit** command related to the cell from the top of the context menu.
- You can delete a rule by selecting it and clicking the **Delete Row** button.
- You can export the entire list of filter rules to a comma-separated values (CSV) file. Click **Export to File**, navigate to an appropriate folder on the Security Manager server, change the file name if you do not like the default name, and click **Save**.

Step 3 Configure the filter rule. Following are the highlights of what you typically need to configure. For specific information on configuring the fields, and for information on fields not mentioned here, see Filter Item Dialog Box, page 40-9.

- Name—You must enter a name for the rule. Use a name that is meaningful to you.
- Signature, Subsignature ID—If the filter should apply to all signatures, use the default values. If you are targeting a specific signature, enter its signature and subsignature identifiers. You can obtain these values by finding the signature in the Signatures policy (see Signatures Page, page 39-4).
- Attacker and Victim Addresses and Ports—If the filter should apply no matter who is attacking, or who is the victim, use the default values. If you are creating a filter specific to an attacker or victim, update these fields to match the appropriate address and port.
- Risk Rating—You are most likely to want to change this value. The filter is applied to events that are within the minimum-maximum range you configure here. The default value, 0-100, will apply the filter rule to all events. If you configure a specific signature ID, the rating applies only to events for that signature (in which case the default risk rating might be acceptable).

  For example, you might want to target only high-risk events, such as 90-100.
- Actions to Subtract—Select the actions that you want to subtract from the event. Use Ctrl+click to select more than one action. If you select an action that is not actually assigned to an event, the filter rule essentially has no effect on the event. For more information about the actions, see Edit, Add, Replace Action Dialog Boxes, page 39-12.

- Stop on Match—Whether to define this filter rule as a stop rule. This setting determines how the remaining rules in the event action filter rules table are processed:
  - If you select this option, and an event meets the conditions of the rule, this rule is the final rule tested for the event. The actions identified by this rule are removed from the event, and the device moves on to perform all remaining actions assigned to the event.
  - If you do not select this option, then events that meet the conditions of this filter rule are also compared to subsequent rules in the event actions filters table. Subsequent rules are tested until either all rules are tested, or the event matches a stop rule.

Click **OK** when you are finished defining your filter rule.

Step 4 If you did not select the right row before adding the rule, select the new rule and use the up and down arrow buttons to position the rule appropriately. Ensure that stop rules are placed after other rules that you want applied prior to the stop.

Tips for Managing Event Action Filter Rules

Following are some tips that might help you effectively manage your event action filter rules:

- Disabled rules are shown with hash marks covering the table row. To change the enabled/disabled status of a rule, right click the rule and select **Enable** or **Disable** as appropriate. You can also change the status when editing the rule.

  Disabling a rule is useful if you want to stop using the rule, but you might want to start using it again in the future. Disabled rules remain in the table so that you do not need to recreate them.

- For existing rules, you can edit most of the fields directly from the event actions filter rules table by right-clicking the cell and selecting the appropriate Edit command from the top portion of the context menu. For example, you can right click the Attacker Ports cell and select **Edit Attacker Ports**.

  Many of these right-click commands open a version of the Edit Filter Item dialog box that contains only the selected property. Other commands simply change a value, or open a sub-menu from which you can select a value to add or remove. For example, right-clicking the Action cell provides four commands:
  - **Add to Actions**—Select from a list of actions to add to those already defined in the rule.
  - **Delete from Actions**—Select from a list of actions defined in the rule to remove from the rule.
  - **Replace Actions With**—Select from a list the action that you want to completely replace those defined in the rule.
  - **Edit Actions**—Opens a dialog box where you can select all actions for the rule. Your selection replaces the cell contents.

- Although filter rules are configured as an ordered list, the rules are not processed as a "first match wins" list, even through they are processed and applied top to bottom. Instead, each rule has a Stop property: the rule is either a stop rule or it is not a stop rule. Processing ends only if an event matches a stop rule. If an event matches a non-stop rule, the event is compared to subsequent filter rules. Thus, more than one filter rule can apply to an event. If you decide to create stop rules, ensure that you place them below all other rules that you want processed for an event.

  If you define no stop rules, each event is compared to all filter rules, and all matching rules are applied to the event in top-to-bottom order.

- You can inherit event action filter rules policies. Thus, you could configure a shared policy in Policy view that includes filter rules that you want to share among all of your devices, inherit that rule for each device (in Device view), and in Device view configure local filter rules that are unique to each device. For more information on inheriting policies, see:
  - Creating a New Shared Policy, page 5-54
  - Inheritance vs. Assignment, page 5-6
  - Inheriting or Uninheriting Rules, page 5-47

Related Topics

- Configuring Event Action Filters, page 40-4
- Event Action Filters Page, page 40-7

Event Action Filters Page

Use the Event Actions Filters page to configure event action filter rules. Filter rules can remove specific actions from an event or they can discard an entire event and prevent further processing by the sensor.

Event action filters are processed as an ordered list and you can move filters up or down in the list. Filters let the sensor perform certain actions in response to the event without requiring the sensor to perform all actions or remove the entire event. Filters work by removing actions from an event. A filter that removes all actions from an event effectively consumes the event.

Before configuring event action filter rules, read the following topics:

- Configuring Event Action Filters, page 40-4
- Tips for Managing Event Action Filter Rules, page 40-6
- Understanding the IPS Event Action Process, page 40-1

Tip Disabled rules are shown with hash marks covering the table row. To change the enabled/disabled status of a rule, right click the rule and select **Enable** or **Disable** as appropriate. You can also change the status when editing the rule.

Navigation Path

- (Device view) Select **IPS > Event Actions > Event Action Filters** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Event Action Filters**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Event Action Filters**, then select an existing policy or create a new one.

Field Reference

Table 40-2 Event Action Filters Page

Element	Description
Name	The name of the filter rule.
Active	Whether the signature is active. This cell is not available for Cisco IOS IPS policies.

Table 40-2 Event Action Filters Page (continued)

Element	Description
IDs	The signature identifiers to which this rule applies.
Subs	The subsignature identifiers.
Attackers	The IP address of the attacker that triggers the filter rule, which can be a host address, an address range (such as 0.0.0.0-255.255.255.255 in the case of IPv4 or ::0-FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF in the case of IPv6), or a network/host policy object. **Tip** If you use a network/host object, you can see the contents of the object by right-clicking it and selecting **Show Contents**. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure.
Attack Ports	The port used by the attacker host that triggers the filter.
Victims	The IP address of the victim that triggers the filter rule, which can be a host address, an address range (such as 0.0.0.0-255.255.255.255 in the case of IPv4 or ::0-FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF in the case of IPv6), or a network/host policy object. **Tip** If you use a network/host object, you can see the contents of the object by right-clicking it and selecting **Show Contents**. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure.
Victim Ports	The port targeted by the attacker host that triggers the filter.
Actions	The actions that should be removed from the event when the filter is triggered.
RR	The risk rating range that triggers this event action filter. For a detailed explanation of how risk rating is calculated, see Calculating the Risk Rating in *Installing and Using Cisco Intrusion Prevention System Device Manager 7.0* on Cisco.com.
Stop	Whether this is a stop rule. If Yes, then when an event meets the conditions of this rule, the filter is applied to the event but the event is not tested against the remaining rules in the event action filter rules policy.
Export to File button	Click this button to export the event action filters summary to a comma-separated values (CSV) file. You are prompted to select the folder on the Security Manager server and to specify a file name.

Table 40-2 Event Action Filters Page (continued)

Element	Description
Up Row and Down Row buttons (arrow icons)	Click these buttons to move the selected rules up or down within a scope. Filter rules are processed in order top to bottom for each event. If the conditions of an event match those defined for a filter, and the filter has the Stop field set to Yes, that filter is applied and no additional filters are considered. Ensure that stop rules are placed after the other rules you want applied to an event. You should order the more restrictive rules before general rules in the table.
Add Row button	Click this button to add a filter rule to the table after the selected row using the Add Filter Item dialog box (see Filter Item Dialog Box, page 40-9). If you do not select a row, the rule is added at the end of the local scope.
Edit Row button	Click this button to edit the selected rule. You can also edit individual cells by right-clicking the cell and selecting the appropriate Edit command.
Delete Row button	Click this button to delete the selected rule. **Tip** Instead of removing the rule, you can right-click the rule and select **Disable**. This prevents the rule from being used, but leaves it in the table in case you want to use it again at a later time.

Filter Item Dialog Box

Use the Add or Edit Filter Item dialog box to configure an event action filter rule.

Tip For existing rules, you can edit most of these fields directly from the event actions filter rules table by right-clicking the cell and selecting the appropriate command from the top portion of the context menu. For example, you can right click the Attacker Ports cell and select **Edit Attacker Ports**. Many of these right-click commands open a version of the Edit Filter Item dialog box that contains only the selected property. When seeking help for these context-editing dialog boxes, look for the property description in the table below.

Navigation Path

From the Event Action Filters page (see Event Action Filters Page, page 40-7), click the **Add Row** button, or select a filter rule and click the **Edit Row** button.

Related Topics

- Configuring Event Action Filters, page 40-4
- Tips for Managing Event Action Filter Rules, page 40-6

Field Reference

Table 40-3 Filter Item Dialog Box

Element	Description
Active Enabled (Active does not apply to Cisco IOS IPS devices.)	Whether the filter rule is active and enabled. Active means that the filter has been put into the filter list and will take effect on filtering events. The default is that the rule is both active and enabled, which means that the rule is used when events are processed. **Tips** • If a filter is active but not enabled, it will still be included in the ordering list; it will be processed, but it will not be used. • If a filter is not active, then it will not be included at all in the ordering of the filters; it will not be processed at all. • Disabled rules are shown in the event action filters table with cross-hatching.
Name	The name of the filter rule. The following characters are allowed in filter names: a-z, A-Z, 0-9, -, . (dot or period), : (colon), and _ (underscore).
Signature IDs	The numerical signature IDs to which the filter rule applies. You can enter a single signature ID, a comma-separated list, or a range of IDs. The default is to apply the rule to signatures in the range 900-65535.
SubSignature ID	The subsignature ID for the specified signature to which the filter rule applies. The subsignature ID identifies a more granular version of a broad signature, but it is not used for all signatures. Enter a subsignature ID appropriate for the signature ID you specified, or enter a range of subsignature IDs. The default value is the range of 0-255.
Attacker IPv4 Address	The IP address of the host that sent the offending packet. You can specify a single host IP address, a range of addresses, or the name of a network/host policy object that identifies the address or address range. Click **Select** to select a network/host object from a list or to create a new object. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure. The default value is a range of all IPv4 addresses (0.0.0.0-255.255.255.255).
Attacker IPv6 Address	The IP address of the host that sent the offending packet. You can specify a single host IP address, a range of addresses, or the name of a network/host policy object that identifies the address or address range. Click **Select** to select a network/host object from a list or to create a new object. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure. The default value is a range of all IPv6 addresses (::0-FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF).

Table 40-3 Filter Item Dialog Box (continued)

Element	Description
Attacker Port	The port used by the attacker host. This is the port from which the offending packet originated. You can also enter a range of ports. The default value is a range of all ports (0-65535).
Victim IPv4 Address	The IP address of the host being attacked (the recipient of the offending packet). You can specify a single host IP address, a range of addresses, or the name of a network/host policy object that identifies the address or address range. Click **Select** to select a network/host object from a list or to create a new object. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure. The default value is a range of all IPv4 addresses (0.0.0.0-255.255.255.255).
Victim IPv6 Address	The IP address of the host being attacked (the recipient of the offending packet). You can specify a single host IP address, a range of addresses, or the name of a network/host policy object that identifies the address or address range. Click **Select** to select a network/host object from a list or to create a new object. **Note** Do not create an IPv4 object and an IPv6 object with the same name; doing so leads to deployment failure. The default value is a range of all IPv6 addresses (::0-FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF).
Victim Port	The port of the host being attacked (the recipient of the offending packet). This is the port to which the offending packet was sent. You can also enter a range of ports. The default value is a range of all ports (0-65535).
Risk Rating Min. and Max.	The risk rating range, between 0 and 100, that should be used to trigger this event action filter. The default value is the complete range (0-100). If an event occurs with a risk rating that falls within the minimum-maximum range you configure here, the event is processed against the rules of this event filter.
OS Relevance	Indicates whether the alert is relevant to the OS that has been identified for the victim. Possible values include one or more of the following: Not Relevant, Relevant, Unknown. Use Ctrl+click to select multiple values. The default is all values selected. **Note** OS Relevance is applicable only to appliances and service modules running IPS 6.x+ software. For Cisco IOS IPS devices, this field is read-only and cannot be edited, and for IPS 5.x devices, this field is blank.
Comments	The user comments associated with this filter, such as an explanation of the purpose of the rule.

Table 40-3 Filter Item Dialog Box (continued)

Element	Description
Actions to Subtract	The actions that should be removed from the event, should the conditions of the event meet the criteria of the event action filter. You can select one or more actions in this list box. All selected actions are removed from the event. Use Ctrl+click to select multiple values. For more information about the possible actions, see Edit, Add, Replace Action Dialog Boxes, page 39-12. For IOS IPS devices, the possible values are restricted to the following: • *Deny Attacker Inline* blocks the attacker's source IP address completely. No connection can be established from the attacker to the router until the shun time expires. You can configure this time in the Event Actions Settings policy as described in Configuring Settings for Event Actions, page 40-23. • *Deny Connection Inline* blocks the appropriate TCP flow from the attacker. Other connections from the attacker can be established to the router. • *Deny Packet Inline* discards the packet without sending a reset. Cisco recommends using "drop and reset" in conjunction with alarm. • *Produce Alert* sends a notification about the attack through syslog or SDEE. • *Reset TCP Connection* is effective for TCP-based connections and sends a reset to both the source and destination addresses. For example, in case of a half-open SYN attack, Cisco IOS IPS can reset the TCP connections.
% to Deny	The percentage of packets to deny for deny attacker features. The range is 0 to 100. The default is 100 percent. **Note** For IOS IPS devices, this field is read only and cannot be edited.
Stop on Match	Whether to define this filter rule as a stop rule. This setting determines how the remaining rules in the event action filter rules table are processed: • If you select this option, and an event meets the conditions of the rule, this rule is the final rule tested for the event. The actions identified by this rule are removed from the event, and the device moves on to perform all remaining actions assigned to the event. • If you do not select this option, then events that meet the conditions of this filter rule are also compared to subsequent rules in the event actions filters table. Subsequent rules are tested until either all rules are tested, or the event matches a stop rule.

Configuring Event Action Overrides

You can add an event action override to change the actions associated with an event based on the risk rating of that event. Event action overrides are a way to add event actions globally without having to configure each signature individually.

Each event action has an associated risk rating range. If a signature event occurs and the risk rating for that event falls within the range for an event action, that action is added to the event. For example, if you want any event with a risk rating of 85 or more to generate an SNMP trap, you can create an event action override for Request SNMP Trap with the risk rating 85-100.

Tip If you want to prevent the use of action overrides, you can disable the entire event action override component as described in Configuring Settings for Event Actions, page 40-23.

Related Topics

- Understanding the IPS Event Action Process, page 40-1

Step 1 Do one of the following to open the Event Action Overrides policy:

- (Device view) Select **IPS > Event Actions > Event Action Overrides** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Event Action Overrides**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Event Action Overrides**, then select an existing policy or create a new one.

The table shows the existing overrides, including the action, the risk rating of the alerts the action will be added to, and whether the rule is enabled. The order of the rules does not matter: all overrides that apply to an alert add the associated actions.

The table can have at most a single entry for each possible action.

Step 2 Configure the desired overrides:

- To add a new override, click the **Add Row (+)** button beneath the table and fill in the Add Event Action Rule dialog box. In the dialog box, select the action you want to add, enter the rating range of the alerts to which you are adding the action (for example, 90-100), and click **OK**. For more information, see Add or Edit Event Action Rule Dialog Box, page 40-14.

  The risk rating range must be between 0 and 100. Separate the low and high of the range with a hyphen, for example, 80-90.

  When adding a new override, you can define your own risk rating, or you can use a pre-defined Risk Rating policy object; beginning with Version 4.5, Security Manager has several pre-defined Risk Rating policy objects:

  - Extreme Risk (90-100)
  - High Risk (76-90)
  - Medium-High Risk (61-75)
  - Medium Risk (46-60)
  - Medium-Low Risk (30-45)
  - Low Risk (16-30)
  - Very Low Risk (1-15)

For more information on these pre-defined policy objects, refer to Configuring Risk Rating Policy Objects, page 40-15.

These pre-defined policy objects cannot be edited, but you can add and edit any of your own policy objects that you have defined.

- To edit an override, to disable it or to change the risk rating, select the override and click the **Edit Row (pencil)** button. You cannot change the event action.

  Note Re-discovery of an IPS device will replace the Risk Rating policy object's value with its inline value. For example, if you assign the High Risk policy object (80-89) to any of the event actions and deploy it to the device, then after re-discovery that policy object's value will be replaced with its inline value of 80-89.

- To remove an override, select it and click the **Delete Row** button.

  Note Policies for IPS appliances and service modules include a default override for Deny Packet Inline, which you cannot delete. If you do not want to use that override, disable it.

- To export the entire list of overrides to a comma-separated values (CSV) file, click **Export to File**, navigate to an appropriate folder on the Security Manager server, change the file name if you do not like the default name, and click **Save**.

Add or Edit Event Action Rule Dialog Box

Use the Add or Edit Event Action Rule dialog box to add an event action rule based on one of the pre-defined Risk Rating policy objects that are available in Security Manager beginning with Version 4.5.

Navigation Path

From the Event Action Overrides policy, click the **Add Row** button beneath the overrides table, or select a row in the table and click the **Edit Row** button. For information on opening the Event Action Overrides policy, see Configuring Event Action Overrides, page 40-13.

Field Reference

Table 40-4 Add or Edit Event Action Rule Dialog Box

Element	Description
Risk Rating	One of several pre-defined Risk Rating policy objects that are available in Security Manager beginning with Version 4.5: • Extreme Risk (90-100) • High Risk (76-90) • Medium-High Risk (61-75) • Medium Risk (46-60) • Medium-Low Risk (30-45) • Low Risk (16-30) • Very Low Risk (1-15) For more information on using one of these pre-defined Risk Rating policy objects, or defining your own, refer to Configuring Event Action Overrides, page 40-13.
Assigned	Whether a particular action is assigned to at least one of the Risk Rating policy objects.
Action Name	The action to be taken for a particular Risk Rating when assigned.
Enabled	Whether a particular action is enabled. Deselect this option to temporarily disable an action without deleting it.

Configuring Risk Rating Policy Objects

Use the Risk Rating Policy Object Dialog Box to configure policy objects for IPS. Seven pre-defined policy objects are available for risk rating; you can also define your own.

Navigation Path

Select **Manage > Policy Objects > All Object Types**, then select **Risk Rating** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**. You cannot edit the pre-defined policy objects, however.

Depending upon whether you selected **New Object** or **Edit Object**, the Add Risk Rating or Edit Risk Rating dialog box appears: refer to Add or Edit Event Action Rule Dialog Box, page 40-14

The remainder of this topic describes the fields that you see in the Risk Rating Policy Object dialog box.

Related Topics

- Configuring Event Action Overrides, page 40-13
- Add or Edit Event Action Rule Dialog Box, page 40-14

Field Reference

Table 40-5 Risk Rating Policy Object Dialog Box

Element	Description
Name	The name of a pre-defined policy object, such as "High Risk," or the name of a policy object that you have defined.
Range	The Risk Rating of a particular policy object, expressed as a numerical range.
Category	Allows you to select Cat-A through Cat-G. This is the category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Overrides	Whether any IPS event action overrides have been configured for this policy object
Description	A text description that you can provide; applies to policy objects that you have defined, not to pre-defined policy objects.
Last Ticket(s)	The last ticket used for this policy object.
Last Modified Date	The last date on which this policy object was modified.

Add or Edit Risk Rating Dialog Box

Use the Add or Edit Risk Rating dialog box to define policy objects for IPS Risk Rating.

Navigation Path

Select **Manage > Policy Objects > All Object Types**, then select **Risk Rating** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**. You cannot edit the pre-defined policy objects, however.

Related Topics

- Configuring Event Action Overrides, page 40-13
- Add or Edit Event Action Rule Dialog Box, page 40-14

Field Reference

Table 40-6 Add or Edit Risk Rating Dialog Box

Element	Description
Name	The name of a pre-defined policy object, such as "High Risk," or the name of a policy object that you have defined.
Description	A text description that you can provide; applies to policy objects that you have defined, not to pre-defined policy objects.
Range	The Risk Rating of a particular policy object, expressed as a numerical range.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 40-6 Add or Edit Risk Rating Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Configuring IPS Event Action Network Information

Use the Event Actions Network Information policy to configure these features:

- **Target value ratings (IPv4 Target Value Ratings tab and IPv6 Target Value Ratings tab)**—You can configure the target value ratings of your network assets. The sensor uses these ratings when calculating the overall risk rating of an alert. By identifying your mission-critical assets, you can trigger more severe signature event actions. As the names indicate, you can use IPv4 or IPv6 by selecting the appropriate tab.

  Target value rating is available on IPS appliances, service modules, and Cisco IOS IPS devices.

  For more information, see Configuring Target Value Ratings, page 40-17.

- **Passive OS fingerprinting and OS mappings (OS Identification tab)**—You can enable the sensor to use information about the operating system running on a device to determine the attack relevance rating, which is a component of the overall risk rating.

  Passive OS fingerprinting and OS mappings are available on devices running IPS 6.x+ software only, and are not available on Cisco IOS IPS devices.

  For more information, see:

  - Understanding Passive OS Fingerprinting, page 40-19
  - Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only), page 40-21

To open the Network Information policy, do one of the following:

- (Device view) Select **IPS > Event Actions > Network Information** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Network Information**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Network Information**, then select an existing policy or create a new one.

Configuring Target Value Ratings

You can assign target value ratings to your network assets. The target value rating is one of the factors used to calculate the risk rating value for each alert. It identifies the perceived importance of a network asset, which you identify by its IP address.

You can develop a security policy that is more stringent for valuable corporate resources and looser for less important resources. For example, you could assign a target value rating to the company web server that is higher than the target value rating you assign to a desktop node. In this example, attacks against the company web server have a higher risk rating than attacks against the desktop node. Events with a higher risk rating trigger more severe signature event actions.

You can configure four value ratings. From highest value to lowest: Mission Critical, High, Medium, Low, No Value (zero value).

For a detailed explanation of how risk rating is calculated, see Calculating the Risk Rating in *Installing and Using Cisco Intrusion Prevention System Device Manager 7.0* on Cisco.com.

Tip If you are configuring target value ratings on a device that uses IPS 6.0 software lower than 6.0(5), you might also want to update the OS Identification tab of the Network Information policy to get around a software bug, even if you do not need to create OS maps. For detailed information, see Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only), page 40-21.

Related Topics

- Configuring IPS Event Action Network Information, page 40-17
- Understanding the IPS Event Action Process, page 40-1

Step 1 Do one of the following to open the Network Information policy:

- (Device view) Select **IPS > Event Actions > Network Information** from the Policy selector, then click the **IPv4 Target Value Ratings** tab or the **IPv6 Target Value Ratings** tab.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Network Information**, then select an existing policy or create a new one. Click the **IPv4 Target Value Ratings** tab or the **IPv6 Target Value Ratings** tab.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Network Information**, then select an existing policy or create a new one. Click the **IPv4 Target Value Ratings** tab.

Note Cisco IOS IPS devices do not support IPv6.

The tab shows the target value ratings that are already configured, showing the IP addresses associated with each configured ratings category. The table can have at most five entries, one per rating category.

Step 2 Configure the desired target value ratings categories:

- To add a new ratings category, click the **Add Row (+)** button beneath the table and fill in the Add Target Value Rating dialog box. In the dialog box, select the rating you want to add, enter the host, network, and address ranges to associate with the category, and click **OK**. For more information, see Target Value Rating Dialog Box, page 40-19.

  For IPv4 addresses, you can specify a single network/host object, or a comma-separated list of host, network, or address ranges, such as 10.10.10.10, 10.10.10.0/24, or 10.10.10.2-10.10.10.254. Addresses that you enter in the network format are converted to address ranges. For IPv6 addresses, use IPv6 addressing conventions.
- To edit the IP addresses in an existing ratings category, select the category and click the **Edit Row (pencil)** button. You cannot change the value rating.

- To remove a rating, select it and click the **Delete Row** button.

Target Value Rating Dialog Box

Use the Add or Edit Target Value Rating dialog box to associate the IP addresses of your assets to a ratings category. IP addresses are IPv4 when you open the Target Value Rating dialog box from the IPv4 Target Value Ratings tab; they are IPv6 from the IPv6 tab.

Navigation Path

From the IPv4 Target Value Rating tab or the IPv6 Target Value Rating tab of the IPS Event Actions Network Information policy, click the **Add Row** button beneath the Target Value Ratings table, or select a row in the table and click the **Edit Row** button. For information on opening the IPv4 Target Value Rating tab or the IPv6 Target Value Rating tab, see Configuring Target Value Ratings, page 40-17.

Field Reference

Table 40-7 Target Value Rating Dialog Box

Element	Description
Value	The target value rating to associate with the specified addresses. From highest to lowest importance: Mission Critical, High, Medium, Low, No Value. This list includes only those value ratings that you have not already configured in the target value ratings table. You change this option when editing a ratings category.
target-address	The IP addresses of the network assets assigned to this value rating. You can specify addresses using the following techniques: • Enter the name of a single network/host object, or click **Select** to select an object from a list or to create a new one. The object can contain a group of networks, hosts, and address ranges. • A comma-separated list of host or network addresses or address ranges. For example, using IPv4, 10.10.10.0/24, 10.10.10.10, 10.10.10.2-10.10.10.254. Addresses that you enter in the network format are converted to address ranges; for example, 10.10.10.0/24 is converted to 10.10.10.0-10.10.10.255.

Understanding Passive OS Fingerprinting

Passive operating system (OS) fingerprinting is enabled by default on IPS 6.0+ sensors and the IPS contains a default vulnerable OS list for each signature.

Passive OS fingerprinting lets the sensor determine the OS that hosts are running. The sensor analyzes network traffic between hosts and stores the OS of these hosts with their IP addresses. The sensor inspects TCP SYN and SYNACK packets exchanged on the network to determine the OS type.

The sensor then uses the OS of the target host OS to determine the relevance of the attack to the victim by computing the attack relevance rating component of the risk rating. Based on the relevance of the attack, the sensor may alter the risk rating of the alert for the attack or the sensor may filter the alert for

the attack. You can then use the risk rating to reduce the number of false positive alerts (a benefit in IDS mode) or definitively drop suspicious packets (a benefit in IPS mode). Passive OS fingerprinting also enhances the alert output by reporting the victim OS, the source of the OS identification, and the relevance to the victim OS in the alert.

Passive OS fingerprinting consists of three components:

- Passive OS learning.

  Passive OS learning occurs as the sensor observes traffic on the network. Based on the characteristics of TCP SYN and SYNACK packets, the sensor makes a determination of the OS running on the host of the source IP address.

- User-configurable OS identification.

  You can configure OS host mappings, which take precedence over learned OS mappings.

- Computation of attack relevance rating and risk rating.

  The sensor uses OS information to determine the relevance of the attack signature to the targeted host. The attack relevance is the attack relevance rating component of the risk rating value for the attack alert.

There are three sources of OS information. The sensor ranks the sources of OS information in the following order:

1. Configured OS mappings—OS mappings that you enter on the OS Identification tab of the Event Actions Network Information policy. You can configure different mappings for each virtual sensor. For more information, see Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only), page 40-21.

   We recommend configuring OS mappings to define the identity of the OS running on critical systems. It is best to configure OS mappings when the OS and IP address of the critical systems are unlikely to change.

2. Imported OS mappings—OS mappings imported from Management Center for Cisco Security Agents (CSA MC).

   Imported OS mappings are global and apply to all virtual sensors. For information on configuring the sensor to use CSA MC, see Configuring the External Product Interface, page 36-26.

3. Learned OS mappings—OS mappings observed by the sensor through the fingerprinting of TCP packets with the SYN control bit set.

   Learned OS mappings are local to the virtual sensor that sees the traffic.

When the sensor needs to determine the OS for a target IP address, it consults the configured OS mappings. If the target IP address is not in the configured OS mappings, the sensor looks in the imported OS mappings. If the target IP address is not in the imported OS mappings, the sensor looks in the learned OS mappings. If it cannot find it there, the sensor treats the OS of the target IP address as unknown.

Tip You can configure Event Action Filter rules to use the OS relevancy value of the target, and configure signatures to identify the OSes vulnerable to a signature.

Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only)

Use the OS Identification tab on the Event Actions Network Information policy to configure operating system (OS) host mappings, which take precedence over learned OS mappings. On the OS Identifications tab you can add, edit, and delete configured OS maps. You can move them up and down in the list to change the order in which the sensor computes the attack relevance rating and risk rating for that particular IP address and OS type combination.

Note OS Identification applies to IPS 6.0+ sensors only and does not apply to Cisco IOS IPS devices.

You can also move them up and down in the list to change the order in which the sensor resolves the OS associated with a particular IP address. Configured OS mappings allow for ranges, so for network 192.168.1.0/24 you might define the following:

IP Address Range Set	OS
192.168.1.1	IOS
192.168.1.2-192.168.1.10,192.168.1.25	UNIX
192.168.1.1-192.168.1.255	Windows

More specific mappings should be at the beginning of the list. Overlap in the IP address range sets is allowed, but the entry closest to the beginning of the list takes precedence.

Tip There is a bug in IPS 6.0 versions lower than 6.0(5) related to the Network Information policy. Even if you change nothing on the OS Identification tab, but you make configuration changes to the Threat Value Ratings tab, Security Manager configures the device to use the **any** variable for restricting OS mappings to addresses. This can result in your monitoring application showing “any” as the event locality for all events. The solution is to upgrade the IPS version on your sensor. The workaround is to enter a non-default value in the **Restrict to these IP Addresses** field on the OS Identification tab, even if you are not configuring specific OS mappings. For example, enter 0.0.0.1-255.255.255.255 instead of “any” or 0.0.0.0-255.255.255.255.

Navigation Path

- (Device view) Select **IPS > Event Actions > Network Information** from the Policy selector, then click the **OS Identification** tab.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Network Information**, then select an existing policy or create a new one. Click the **OS Identification** tab.

Related Topics

- Configuring IPS Event Action Network Information, page 40-17
- Understanding the IPS Event Action Process, page 40-1

Field Reference

Table 40-8 OS Identification Tab

Element	Description
Enable Passive OS Fingerprinting	When selected, lets the sensor perform passive OS analysis. You must enable this option for any of the maps configured on this page to be used. Passive OS fingerprinting functions as part of the sensor. As the sensor analyzes network traffic between hosts, the sensor stores the identity of the OS running on the hosts alongside the IP addresses of the hosts. The sensor determines the identity of the OSes on the hosts by inspecting characteristics of the packets exchanged on the network. The sensor then uses the target system's OS information to compute the ARR (Attack Relevance Rating) component for the RR (Risk Rating). The RR can then be used to drop suspicious packets. For more information about passive OS fingerprinting, see Understanding Passive OS Fingerprinting, page 40-19.
Restricted to these IP Addresses	Restricts attack relevance rating calculation to the specified addresses. You can specify addresses using the following techniques: • Enter the name of a single network/host object, or click **Select** to select an object from a list or to create a new one. The object can contain a group of networks, hosts, and address ranges. • A comma-separated list of host or network addresses or address ranges. For example, 10.10.10.0/24, 10.10.10.10, 10.10.10.2-10.10.10.254.
OS Maps table	The list of OS mappings, showing the IP addresses of the hosts and the operating systems to which they are mapped. When looking for a match, the sensor goes from top to bottom and selects the first rule that matches the IP address. • To add a mapping, click the **Add Row** button and fill in the Add OS Map dialog box (see OS Map Dialog Box, page 40-22). • To edit a mapping, select the rule and click the **Edit Row** button. • To delete a map, select it and click the **Delete Row** button. • To change the priority of a rule, select it and click the **Up** or **Down** arrow buttons until the rule is positioned correctly.

OS Map Dialog Box

Use the Add or Edit OS Map dialog box to map a host through its IP address to an OS type. Create mappings only if you want to statically assign an OS type to an IP address. Because the sensor uses passive OS fingerprinting to discover the OS associated with an IP address, you might not want to create any mappings, or create mappings only for mission-critical devices that have static IP addresses. Update any mappings that you create if you install devices with different operating systems on the address.

Navigation Path

From the OS Identification tab of the IPS Event Actions Network Information policy, click the **Add Row** button beneath the OS Maps table, or select a row in the table and click the **Edit Row** button. For information on opening the OS Identification tab, see Configuring OS Identification (Cisco IPS 6.x and Later Sensors Only), page 40-21.

Field Reference

Table 40-9 OS Map Dialog Box

Element	Description
IP Addresses	The IP addresses for this mapping. You can specify addresses using the following techniques: • Enter the name of a single network/host object, or click **Select** to select an object from a list or to create a new one. The object can contain a group of networks, hosts, and address ranges. • A comma-separated list of host or network addresses or address ranges. For example, 10.10.10.0/24, 10.10.10.10, 10.10.10.2-10.10.10.254.
OS Type	The operating system running on the identified hosts. Select the most appropriate option from the list. You can select multiple options (using Ctrl+click) to indicate that there is more than one possible OS. **Tip** Because these mappings take precedence over learned mappings, you probably are better off not assigning General OS, Other, or Unknown OS. The sensor might be able to learn the actual OS through passive OS fingerprinting and provide a better matching. For more information, see Understanding Passive OS Fingerprinting, page 40-19.

Configuring Settings for Event Actions

Use the Event Actions Settings policy to configure general settings that apply globally to event action rules. The defaults for these options are appropriate for most situations, so change them only if you are certain that your situation requires non-default behavior.

To configure the Event Actions Settings policy, do one of the following:

- (Device view) Select **IPS > Event Actions > Settings** from the Policy selector.
- (Policy view, IPS appliances and service modules) Select **IPS > Event Actions > Settings**, then select an existing policy or create a new one.
- (Policy view, Cisco IOS IPS devices) Select **IPS (Router) > Event Actions > Event Action Settings**, then select an existing policy or create a new one.

The following table describes the options you can configure. Note that the options available for Cisco IOS IPS devices are more limited than those available for IPS appliances and service modules.

Tip Do not disable the Summarizer except for troubleshooting purposes. If you disable the Summarizer, every signature is set to Fire All with no summarization. Note that you do not need to change the state of the Meta Event Generator. Cisco has discontinued the use of Meta signatures, and they have all been retired.

Table 40-10 Event Actions Settings Policy

Element	Description
Enable Event Action Override (All device types.)	When selected, enables override rules as defined on the Event Action Overrides page. You can add an event action override to add actions to an event based on specific details about that event. For configuring override rules, see Configuring Event Action Overrides, page 40-13.
Enable Event Action Filters (All device types.)	When selected, enables the filter rules as defined on the Event Action Filters page. You can configure event action filters to remove specific actions from an event or to discard an entire event and prevent further processing by the sensor. For configuring event action filters rules, see Configuring Event Action Filters, page 40-4.
Enable Event Action Summarizer (IPS appliances and service modules only.)	When selected, enables the Summarizer component. The Summarizer groups events into a single alert, thus decreasing the number of alerts the sensor sends out. By default, the Summarizer is enabled. If you disable it, all signatures are set to Fire All with no summarization. If you configure individual signatures to summarize, this configuration is ignored when the Summarizer is not enabled. The Report Manager component of Cisco Security Manager reports events individually. The Event Viewer component of Cisco Security Manager displays alerts. As stated above, the Summarizer groups events into a single alert, thus decreasing the number of alerts the sensor sends out. **Tip** Cisco IPS Manager Express (IME) and Cisco Security Manager do not summarize events in precisely the same way.
Enable Meta Event Generator (IPS appliances and service modules only.)	Cisco recommends that you do not change the state of the Meta Event Generator. Cisco has discontinued the use of Meta signatures, and they have all been retired.

Table 40-10 Event Actions Settings Policy (continued)

Element	Description
Enable Threat Rating Adjustment (IPS appliances and service modules only.)	When selected, enables threat rating adjustment, which adjusts the risk rating. If disabled, risk rating is equal to threat rating. Available in sensors running IPS 6.0+ software only. The Threat Rating feature provides a single view of the threat environment of the network. Threat Rating minimizes alarms and events through a customized view that shows only events with a high Threat Rating value. The Threat Rating value is derived as follows: • Dynamic adjustment of event Risk Rating based on success of response action • If response action was applied, Risk Rating is deprecated (Threat Rating < Risk Rating) • If response action was not applied, Risk Rating remains unchanged (Threat Rating = Risk Rating) The result is a single value by which the threat risk is determined.
Deny Attacker Duration in seconds (All device types.)	The number of seconds to deny the attacker inline. The range is 0 to 518400. The default is 3600.
Block Attack Duration in minutes (IPS appliances and service modules only.)	The number of minutes to block a host or connection. The range is 0 to 10000000. The default is 30.
Maximum Number of Denied Attackers (IPS appliances and service modules only.)	Limits the number of denied attackers possible in the system at any one time. The range is 0 to 100000000. The default is 10000.
Enable One Way TCP Reset (IPS appliances and service modules only.)	When selected, enables a one-way TCP reset for deny packet inline actions for TCP-based alerts. Available only for sensors running IPS 6.1+ software. The one-way TCP reset operates for inline mode only and is an automatic addition to the deny packet inline actions. It sends a TCP reset to the victim of the alert, thus creating a black hole for the attacker and clearing the TCP resources of the victim. **Tips** • In inline mode, all packets entering or leaving the network must pass through the sensor. • An inline sensor denies packets for any alert with a risk rating of greater than or equal to 90. It also issues a one-way TCP reset on TCP alerts with a risk rating of greater than or equal to 90.

Managing IPS Anomaly Detection

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

Anomaly detection is designed to recognize network congestion caused by worm traffic that exhibits scanning behavior. Anomaly detection also will identify infected hosts on the network that are scanning for other vulnerable hosts.

Anomaly detection is enabled by default, but there are some configuration settings you should adjust to use it effectively.

Note The sensor must use IPS software version 6.x or higher to configure anomaly detection. In addition, Cisco IOS IPS and the AIP-SSC-5 do not support anomaly detection.

This chapter contains the following topics:

- Understanding Anomaly Detection, page 41-1
- Configuring Anomaly Detection, page 41-6

Understanding Anomaly Detection

The anomaly detection component of the sensor detects worm-infected hosts. This enables the sensor to be less dependent on signature updates for protection again worms and scanners, such as Code Red and SQL Slammer and so forth. The anomaly detection component lets the sensor learn normal activity and send alerts or take dynamic response actions for behavior that deviates from what it has learned as normal behavior.

Note Anomaly detection does not detect email-based worms, such as Nimda.

Anomaly detection detects the following two situations:

- When the network starts on the path of becoming congested by worm traffic.
- When a single worm-infected source enters the network and starts scanning for other vulnerable hosts.

The following topics explain anomaly detection in more detail:

- Worm Viruses, page 41-2
- Anomaly Detection Modes, page 41-2
- Anomaly Detection Zones, page 41-3
- Knowing When to Turn Off Anomaly Detection, page 41-4
- Configuring Anomaly Detection Signatures, page 41-4
- Configuring Anomaly Detection, page 41-6

Worm Viruses

Worm viruses are automated, self-propagating, intrusion agents that make copies of themselves and then facilitate their spread. Worm viruses attack a vulnerable host, infect it, and then use it as a base to attack other vulnerable hosts. They search for other hosts by using a form of network inspection, typically a scan, and then propagate to the next target. A scanning worm virus locates vulnerable hosts by generating a list of IP addresses to probe, and then contacts the hosts. Code Red worm, Sasser worm, Blaster worm, and the Slammer worm are examples of worms that spread in this manner.

Anomaly detection identifies worm-infected hosts by their behavior as a scanner. To spread, a worm virus must find new hosts. It finds them by scanning the Internet using TCP, UDP, and other protocols to generate unsuccessful attempts to access different destination IP addresses. A scanner is defined as a source IP address that generates events on the same destination port (in TCP and UDP) for too many destination IP addresses.

The events that are important for TCP are non-established connections, such as a SYN packet that does not have its SYN-ACK response for a given amount of time. A worm-infected host that scans using TCP generates non-established connections on the same destination port for an anomalous number of IP addresses.

The events that are important for UDP are unidirectional connections, such as a UDP connection where all packets are going in only one direction. A worm-infected host that scans using UDP generates UDP packets but does not receive UDP packets on the same IP address within a time-out period on the same destination port for multiple destination IP addresses.

The events that are important for other protocols, such as ICMP (protocol number 1), are from a source IP address to many different destination IP addresses, that is, packets that are received in only one direction.

Caution If a worm virus has a list of IP addresses it should infect and does not have to use scanning to spread itself (for example, it uses passive mapping—listening to the network as opposed to active scanning), it will not be detected by anomaly detection worm policies. Worm viruses that receive a mailing list from probing files within the infected host and email this list will not be detected, because no Layer 3 or Layer 4 anomaly is generated.

Anomaly Detection Modes

Anomaly detection initially conducts a “peacetime” learning process when the most normal state of the network is reflected. Anomaly detection then derives a set of policy thresholds that best fit the normal network. This is done in two phases: an initial learning mode phase, followed by the ongoing operational detect mode phase.

Anomaly detection has the following modes:

- Learning accept mode (initial setup)

  Although anomaly detection is in detect mode by default, it conducts an initial learning accept mode for the default period of 24 hours. We assume that during this phase no attack is being carried out. Anomaly detection creates an initial baseline, known as a knowledge base, of the network traffic. The default interval value for periodic schedules is 24 hours and the default action is rotate, meaning that a new knowledge base is saved and loaded, and then replaces the initial knowledge base after 24 hours.

  Keep the following in mind:

  - Anomaly detection does not detect attacks when working with the initial knowledge base, which is empty. After the default of 24 hours, a knowledge base is saved and loaded and now anomaly detection also detects attacks.
  - Depending on your network complexity, you may want to have anomaly detection in learning accept mode for longer than the default 24 hours. You configure the mode in the Virtual Sensors policy; see Defining A Virtual Sensor, page 38-5. After your learning period has finished, edit the virtual sensor and change the mode to Detect.

- Detect mode

  For ongoing operation, the sensor should remain in detect mode. This is for 24 hours a day, 7 days a week. Once a knowledge base is created and replaces the initial knowledge base, anomaly detection detects attacks based on it. It looks at the network traffic flows that violate thresholds in the knowledge base and sends alerts. As anomaly detection looks for anomalies, it also records gradual changes to the knowledge base that do not violate the thresholds and thus creates a new knowledge base. The new knowledge base is periodically saved and takes the place of the old one thus maintaining an up-to-date knowledge base.

- Inactive mode

  You can turn anomaly detection off by putting it in inactive mode. Under certain circumstances, anomaly detection should be in inactive mode, for example, if the sensor is running in an asymmetric environment. Because anomaly detection assumes it gets traffic from both directions, if the sensor is configured to see only one direction of traffic, anomaly detection identifies all traffic as having incomplete connections, that is, as scanners, and sends alerts for all traffic flows.

The following example summarizes the default anomaly detection configuration. If you add a virtual sensor at 11:00 pm and do not change the default anomaly detection configuration, anomaly detection begins working with the initial knowledge base and only performs learning. Although it is in detect mode, it cannot detect attacks until it has gathered information for 24 hours and replaced the initial knowledge base. At the first start time (10:00 am by default), and the first interval (24 hours by default), the learning results are saved to a new knowledge base and this knowledge base is loaded and replaces the initial knowledge base. Because the anomaly detection is in detect mode by default, now that anomaly detection has a new knowledge base, the anomaly detection begins to detect attacks.

Anomaly Detection Zones

By subdividing the network into zones, you can achieve a lower false negative rate. A zone is a set of destination IP addresses. There are three zones, each with its own thresholds: internal, illegal, and external.

The external zone is the default zone with the default Internet range of 0.0.0.0-255.255.255.255. By default, the internal and illegal zones contain no IP addresses. Packets that do not match the set of IP addresses in the internal or illegal zone are handled by the external zone.

We recommend that you configure the internal zone with the IP address range of your internal network. If you configure it in this way, the internal zone is all the traffic that comes to your IP address range, and the external zone is all the traffic that goes to the Internet.

You can configure the illegal zone with IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses or part of your internal IP address range that is unoccupied. An illegal zone can be very helpful for accurate detection, because we do not expect any legal traffic to reach this zone. This allows very low thresholds, which in turn can lead to very quick worm virus detection.

Knowing When to Turn Off Anomaly Detection

Anomaly detection assumes that it gets traffic from both directions. If the sensor is configured to see only one direction of traffic, you should turn off anomaly detection. Otherwise, when anomaly detection is running in an asymmetric environment, it identifies all traffic as having incomplete connections, that is, as scanners, and sends alerts for all traffic flows.

You turn off anomaly detection in the Virtual Sensors policy. Edit the virtual sensor for which you are disabling anomaly detection, and change the Anomaly Detection Mode to Inactive. For information on editing virtual sensors, see Editing Policies for a Virtual Sensor, page 38-9.

Configuring Anomaly Detection Signatures

The Traffic Anomaly engine contains nine anomaly detection signatures covering three protocols (TCP, UDP, and other). Each signature has two subsignatures, one for the scanner and the other for the worm-infected host (or a scanner under worm attack). When anomaly detection discovers an anomaly, it triggers an alert for these signatures. All anomaly detection signatures are enabled by default and the alert severity for each one is set to high.

When a scanner is detected but no histogram anomaly occurred, the scanner signature fires for that attacker (scanner) IP address. If the histogram signature is triggered, the attacker addresses that are doing the scanning each trigger the worm signature (instead of the scanner signature). The alert details state which threshold is being used for the worm detection now that the histogram has been triggered. From that point on, all scanners are detected as worm-infected hosts.

The following anomaly detection event actions are possible:

- Produce alert—Writes the event to the Event Store.
- Deny attacker inline—(Inline only) Does not transmit this packet and future packets originating from the attacker address for a specified period of time.
- Log attacker packets—Starts IP logging for packets that contain the attacker address.
- Deny attacker service pair inline—Blocks the source IP address and the destination port.
- Request SNMP trap—Sends a trap notification to an SNMP trap destination. To use this action, you must configure SNMP trap hosts as described in Configuring SNMP, page 36-8.
- Request block host—Sends a request to ARC to block this host (the attacker). To use this action, you must configure blocking devices as described in Configuring IPS Blocking and Rate Limiting, page 43-7.

You can add actions to the signatures either directly, in the Signatures policy, or to events generated by the signatures based on risk rating in the Event Actions Overrides policy.

The following table lists the anomaly detection worm signatures.

Table 41-1 Anomaly Detection Worm Signatures

Signature ID	Subsignature ID	Name	Description
13000	0	Internal TCP Scanner	Identified a single scanner over a TCP protocol in the internal zone.
13000	1	Internal TCP Scanner	Identified a worm attack over a TCP protocol in the internal zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13001	0	Internal UDP Scanner	Identified a single scanner over a UDP protocol in the internal zone.
13001	1	Internal UDP Scanner	Identified a worm attack over a UDP protocol in the internal zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13002	0	Internal Other Scanner	Identified a single scanner over an Other protocol in the internal zone.
13002	1	Internal Other Scanner	Identified a worm attack over an Other protocol in the internal zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.
13003	0	External TCP Scanner	Identified a single scanner over a TCP protocol in the external zone.
13003	1	External TCP Scanner	Identified a worm attack over a TCP protocol in the external zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13004	0	External UDP Scanner	Identified a single scanner over a UDP protocol in the external zone.
13004	1	External UDP Scanner	Identified a worm attack over a UDP protocol in the external zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13005	0	External Other Scanner	Identified a single scanner over an Other protocol in the external zone.
13005	1	External Other Scanner	Identified a worm attack over an Other protocol in the external zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.
13006	0	Illegal TCP Scanner	Identified a single scanner over a TCP protocol in the illegal zone.

Table 41-1 Anomaly Detection Worm Signatures (continued)

Signature ID	Subsignature ID	Name	Description
13006	1	Illegal TCP Scanner	Identified a worm attack over a TCP protocol in the illegal zone; the TCP histogram threshold was crossed and a scanner over a TCP protocol was identified.
13007	0	Illegal UDP Scanner	Identified a single scanner over a UDP protocol in the illegal zone.
13007	1	Illegal UDP Scanner	Identified a worm attack over a UDP protocol in the illegal zone; the UDP histogram threshold was crossed and a scanner over a UDP protocol was identified.
13008	0	Illegal Other Scanner	Identified a single scanner over an Other protocol in the illegal zone.
13008	1	Illegal Other Scanner	Identified a worm attack over an Other protocol in the illegal zone; the Other histogram threshold was crossed and a scanner over an Other protocol was identified.

Configuring Anomaly Detection

Use the Anomaly Detection policy to configure anomaly detection settings. The Virtual Sensors policy also contains a setting important for anomaly detection.

This procedure explains the overall configuration of anomaly detection. Before configuring these settings, read the following topics:

- Understanding Anomaly Detection, page 41-1
- Worm Viruses, page 41-2
- Anomaly Detection Modes, page 41-2
- Anomaly Detection Zones, page 41-3
- Knowing When to Turn Off Anomaly Detection, page 41-4
- Configuring Anomaly Detection Signatures, page 41-4

Step 1 Do one of the following to open the Anomaly Detection policy you want to modify:

- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
- (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.

The Anomaly detection policy includes these tabs:

- **Operation Settings**—Defines the worm timeout and identifies any IP addresses that should be ignored by anomaly detection.

- **Learning Accept Mode**—The configuration for learning mode, including how the knowledge base is handled.
- **Internal Zone, Illegal Zone, External Zone**—The zones of your network that you define. You can configure unique settings for each zone. For an explanation of the zones, see Anomaly Detection Zones, page 41-3.

Step 2 Click the **Operation Settings** tab, if necessary, and configure the following:

- **Worm Timeout**—The time in seconds for the worm termination timeout. The range is 120 to 10,000,000 seconds. The default is 600 seconds. For an explanation of how this timeout is used, see Understanding Anomaly Detection Thresholds and Histograms, page 41-9.
- **Enable Ignored Addresses** and **Source/Destination Addresses to Ignore**—Whether you are configuring a list of addresses that should be ignored while anomaly detection is processing. You can specify a list of source addresses (those that initiate a scan) or destination addresses (the hosts that are scanned).

  The addresses can be single host (such as 10.100.10.1), a range of addresses (such as 10.100.10.0-10.100.10.255), or network/host objects that contain single hosts, address ranges, or a combination of hosts and ranges. Click **Select** to select objects from a list or to create new objects.

Step 3 Click the **Learning Accept Mode** tab and define how the knowledge base will be generated and used. For detailed information, see Configuring Anomaly Detection Learning Accept Mode, page 41-8.

Step 4 Configure the internal, illegal, and external zones:

- Define the internal and illegal zones—The internal zones are the IP addresses of your internal network, the network that you manage. The illegal zone should represent IP address ranges that should never be seen in normal traffic, for example, unallocated IP addresses or part of your internal IP address range that is unoccupied.

  Click the **Internal Zone** and **Illegal Zone** tabs in turn and configure the following on the **General** tab:

  - **Enable this zone**—Whether the zone will be processed by anomaly detection.
  - **Service Subnets**—The IP addresses that comprise the zone. The default (0.0.0.0) is that no address is included in the zone. Replace 0.0.0.0 to define addresses for the zone.

    The addresses can be single host (such as 10.100.10.1), a range of addresses (such as 10.100.10.0-10.100.10.255), or network/host objects that contain single hosts, address ranges, or a combination of hosts and ranges. Click **Select** to select objects from a list or to create new objects.
- Decide whether to enable the external zone—The external zone comprises all IP addresses that are not configured for the internal or illegal zones. You do not explicitly assign addresses to this zone. On the **External Zone** tab, **General** sub-tab, you can enable or disable the zone using the **Enable this zone** checkbox. The external zone is enabled by default.
- Configure scanner thresholds and histograms—Each zone has sub-tabs for **TCP Protocol**, **UDP Protocol**, and **Other Protocols**. On these tabs, you can configure non-default settings for specific services that override the learned histograms. For detailed information about configuring these settings, see Configuring Anomaly Detection Thresholds and Histograms, page 41-11.

At this point, you have finished configuring the basic anomaly detection settings.

Step 5 (Device view only.) Configure the anomaly detection mode. This setting is defined in the **Virtual Sensors** policy. Consider the following tips to select the correct policy:

- If you configured the anomaly detection policy on a virtual sensor (other than vs0, which is represented by the parent IPS device), you must select the parent IPS device, then select the Virtual Sensors policy.

- If you configured the Anomaly Detection policy as a shared policy in Policy view, select the IPS device to which the policy is assigned, or that hosts a virtual sensor to which the policy is assigned.

Then, complete the following steps in the Virtual Sensors policy:

a. Select the desired virtual sensor in the table and click the **Edit Row** button.

b. In the Modify Virtual Sensors dialog box, select the appropriate option for the Anomaly Detection Mode setting: Detect, Inactive, Learn. The default and normal operational mode is Detect. However, if you are using asymmetric normalizer mode, you might want to set the anomaly detection mode to Inactive. For detailed information about these modes, see Anomaly Detection Modes, page 41-2. For information about the other settings in this dialog box, see Virtual Sensor Dialog Box, page 38-7.

c. If you placed anomaly detection in Learning mode, remember to change the mode to Detect after the desired learning period has completed.

Step 6 Add additional actions to the anomaly detection signatures, if desired. For example, you might want to add a deny action so that attacks are dropped. You can alternatively configure event action overrides to add actions based on risk rating. For more information, see Configuring Anomaly Detection Signatures, page 41-4.

Step 7 Manage the knowledge base, if necessary.

If you configured the knowledge base to automatically rotate (on the Learning Accept Mode tab), then the knowledge base is refreshed automatically and manual intervention is not necessary.

If you configured anomaly detection to only save new databases, and not use them, then you need to manually load updated knowledge bases periodically. You cannot do this in Security Manager; use the IPS Device Manager (IDM) instead.

Using IDM (or IME), you can load, delete, and rename knowledge bases, and upload them to or download them from an external server. For more information about what you can do, see the online help for IDM or IME.

Configuring Anomaly Detection Learning Accept Mode

Use the Learning Accept Mode tab of the Anomaly Detection policy to configure whether you want the sensor to create a new knowledge base every so many hours. You can configure whether the knowledge base is created and loaded (Rotate) or saved (Save Only). You can schedule how often and when the knowledge base is loaded or saved.

The default generated filename is YYYY-Mon-dd-hh_mm_ss (that is, year-month-day-hour_minute_second),where Mon is a three-letter abbreviation of the current month.

The knowledge base has a tree structure and contains the following information:

- Knowledge base name
- Zone name
- Protocol
- Service

The knowledge base holds a scanner threshold and a histogram for each service. If you have learning accept mode set to automatic and the action set to rotate, a new knowledge base is created every 24 hours and used in the next 24 hours. If you have learning accept mode set to automatic and the action is set to save only, a new knowledge base is created but not loaded, and the current knowledge base is used. If you do not have learning accept mode set to automatic, no knowledge base is created.

Tip Although you can use Security Manager to configure how knowledge bases are generated, you cannot manage the knowledge bases themselves. Use the IPS Device Manager (IDM), or IPS Manager Express (IME) instead. Using IDM (or IME), you can load, delete, and rename knowledge bases, and upload them to or download them from an external server. For more information about what you can do, see the online help for IDM or IME.

Related Topics

- Anomaly Detection Modes, page 41-2
- Configuring Anomaly Detection, page 41-6
- Understanding Anomaly Detection Thresholds and Histograms, page 41-9

Step 1 Do one of the following to open the Anomaly Detection policy you want to modify:

- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
- (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.

Step 2 Click the **Learning Accept Mode** tab and configure the following options:

- **Automatically accept learning knowledge base**—Whether to have the sensor automatically update the knowledge base. If you do not select this option, anomaly detection does not automatically create a new knowledge base, and you cannot configure the other options on this tab.
- **Action**—Whether to rotate or save the knowledge base when it is created.

  If you choose **Rotate** (the default), the new knowledge base is created and loaded according to the schedule you define. If you choose **Save Only**, the new knowledge base is created but not loaded. You can examine it and decide whether to load it into anomaly detection using IDM or IME.

Step 3 In the **Schedule** field, select the schedule for generating a new knowledge base. The default schedule is periodic starting at 10 AM and running for 24 hours. Options are:

- **Periodic**—Base the schedule on a recurring period. Configure the following options:
  - **Start Time**—The starting time for the learning window in hh:mm:ss format (24-hour clock).
  - **Learning Interval in hours**—How long you want anomaly detection to learn from the network before creating a new knowledge base.
- **Calendar Schedule**—Base the schedule on specific times of day and days of the week. The dialog box changes to show Time of Day and Days of the Week tables. These times apply to every day selected; you cannot specify different times for different days.
  - To add a time or day, click the **Add Row (+)** button beneath the appropriate table. The time is in hh:mm:ss format (24-hour clock). For day, select the day from the list.
  - To edit an existing time or day, select it and click the **Edit Row (pencil)** button.
  - To delete a time or day, select it and click the **Delete Row (trash can)** button. Ensure that you have at least one time and one day configured.

Understanding Anomaly Detection Thresholds and Histograms

Anomaly detection uses thresholds and histograms to determine if scanning behavior is an attack.

During learning mode, anomaly detection develops histograms for each TCP and UDP port, and for other protocols, to create a baseline of the normal behavior of your network (see Anomaly Detection Modes, page 41-2). For example, the histogram for a TCP port lists the "normal" number of source addresses that make incomplete connections to a certain number of destination addresses during a minute. The histograms contain three buckets: low number of destination addresses (5), medium number (20), and high number (100). (The destination buckets are a fixed number.) Separate histograms are kept for each service and zone (see Anomaly Detection Zones, page 41-3).

For example, learning mode might develop the following histogram for TCP port 80:

Number of Destination Addresses	Number of Source Addresses
Low (5)	18
Medium (20)	6
High (100)	2

In addition to these learned histograms, the anomaly detection scanner has a threshold setting that you configure. You configure a general scanner threshold, and you can override the threshold (configuring a different value) for any specific service (TCP port, UDP port, or other protocol). Each zone has its own thresholds.

When anomaly detection moves to detect mode, where it is actively scanning for worms, the thresholds and histograms are used as follows:

- Histograms are ignored until the threshold for the service is exceeded. For example, consider the above table for TCP/80 traffic. If the threshold is set at 200 (the default), a scanner must scan 200 hosts in a minute to trigger a scanner alert. If 7 source addresses scanned 50 hosts (which is an anomaly in the histogram, which expects no more than 6 hosts scanning 20-99 destinations), but a single scanner scanned only 100 addresses, no alert is generated, and no anomaly is detected.
- When the scanner threshold is exceeded, anomaly detection uses the histogram to determine if the service is under worm attack. In this example, if a source scans more than 200 destinations, anomaly detection evaluates the collected activity in the network. Because 7 hosts have scanned 50 hosts, a worm alert is generated.

  When under worm attack, anomaly detection stops learning and clears the current learning information. It also temporarily lowers the thresholds.
- When a worm attack is detected, the worm timeout counter is started. When the timeout is reached, the scanner is reset. If the worm attack continues, new alerts are generated. You configure the worm timeout on the Operation Settings tab of the Anomaly Detection policy.

If you leave the defaults in place, anomaly detection generates histograms based on what it learns about your network from the network's actual behavior. However, if you understand your network you can fine-tune these histograms to reduce false positives, creating your own definition of expected (or desired or tolerated) behavior for each TCP/UDP port or other protocol, for each zone. You can create your own histograms for only those services that interest you, and leave the defaults for all other ports. Additionally, you can configure the general scanner threshold for each zone, and configure different thresholds for specific services.

For information on configuring thresholds and histograms, see Configuring Anomaly Detection Thresholds and Histograms, page 41-11.

Configuring Anomaly Detection Thresholds and Histograms

Anomaly detection uses thresholds and histograms to determine if scanning behavior is an attack. In most cases, you can use the default thresholds and the histograms that anomaly detection generates during learning mode (see Anomaly Detection Modes, page 41-2). However, you might want to fine-tune these settings. Changing the thresholds is a more likely change than creating your own histograms.

Before you configure these settings, read Understanding Anomaly Detection Thresholds and Histograms, page 41-9. You must understand how thresholds and histograms are used together to configure them.

Step 1 Do one of the following to open the Anomaly Detection policy you want to modify:

- (Device view) Select **IPS > Anomaly Detection** from the Policy selector.
- (Policy view) Select **IPS > Anomaly Detection** from the Policy selector. Select an existing policy or create a new one.

Step 2 Click the tab for the zone whose thresholds or histograms you want to change. You configure separate values for each zone: **Internal Zone**, **Illegal Zone**, **External Zone**. For an explanation of the zones, see Anomaly Detection Zones, page 41-3.

The tabs for each zone contain four sub-tabs: General, TCP Protocol, UDP Protocol, and Other Protocols. The General tab defines the IP addresses for the zone and whether the zone is enabled (the External zone includes all IP addresses not specified for the other zones, you do not configure specific addresses for the External zone).

The other tabs are where you define thresholds and histograms.

Step 3 Select the tab for the protocol for which you want to modify thresholds or histograms: **TCP Protocol**, **UDP Protocol**, **Other Protocol**.

On each tab, configure the following options:

- **Enabled**—Whether anomaly detection is enabled for the protocol. You can turn off detection for all of TCP, UDP, or for all non-TCP/UDP protocols with this option. If you deselect the option, any other settings configured on the tab are ignored.
- **Destination Port Map** or **Protocol Number Map** table—This table lists the TCP/UDP ports, or other protocols, for which you are configuring non-default mappings. By default, all ports and protocols are enabled and use the default scanner threshold.

  Add items to this table only if you want to: disable detection for a port or protocol; set a different threshold value for a port or protocol; or configure an explicit histogram for a port or protocol, which will be used instead of the learned histogram.

  - To add a mapping, click the **Add Row (+)** button and fill in the Add Dest or Protocol Map dialog box. For detailed information, see Dest Port or Protocol Map Dialog Box, page 41-12.
  - To edit a mapping, select it and click the **Edit Row (pencil)** button.
  - To delete a mapping, select it and click the **Delete Row (trash can)** button. Deleting a mapping returns the service to the default settings.
- **Scanner Threshold**—The threshold for all TCP, UDP, or other protocols. This threshold is used for all services except those for which you configured a scanner override in the mapping table. The range is 5 to 1000. The default is 200.
- **Threshold Histogram**—The default histogram for all TCP, UDP, or other protocols. This histogram is used for all services except those for which you configured a scanner override in the mapping table.

The content of this table is fixed; you cannot add or delete items. However, you can select a row and click **Edit Row (pencil)** to change the number of source addresses configured for a threshold setting. See Histogram Dialog Box, page 41-13.

Step 4 Repeat the process for each combination of zone and protocol for which you are defining non-default settings.

Dest Port or Protocol Map Dialog Box

Use the Add or Modify Dest Port Map dialog box to add or modify destination port scanner settings for TCP or UDP, and the Add or Modify Protocol Map dialog box to add or modify scanner settings for other protocols.

Before you configure these settings, read the following topics:

- Understanding Anomaly Detection Thresholds and Histograms, page 41-9
- Configuring Anomaly Detection Thresholds and Histograms, page 41-11

Tip You do not need to add a port or protocol to have anomaly detection look for worm attacks against it. By default, all ports and protocols are processed. You need to configure specific settings only if you want to turn off detection on a specific port or protocol, or if you want non-default thresholds or histograms.

Navigation Path

In the Anomaly Detection policy, on the **TCP Protocol**, **UDP Protocol**, or **Other Protocol** sub tabs on the Internal Zone, Illegal Zone, or External Zone tabs, click the **Add Row** button beneath the Destination Port Map or Protocol Number Map tables, or select a row and click the **Edit Row** button. For more information on the steps required to get here, see Configuring Anomaly Detection Thresholds and Histograms, page 41-11.

Field Reference

Table 41-2 Destination Port or Protocol Map Dialog Box

Element	Description
Destination Port Number (Dest Port Map dialog box only.)	The destination port number for which you are defining non-default values. The range is 0 to 65535. Enter a single port number, or the name of a port list object that contains a single port number. Click **Select** to select an object from a list or to create a new one.
Protocol Number (Protocol Map dialog box only.)	The protocol number for non-TCP/UDP protocols. For a list of protocol numbers, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "Protocol Numbers." Look for a heading (at the time of this writing, the second search hit). The range is 0 to 255. For example, ICMP is protocol 1.
Enabled	Whether to enable this service. If you do not enable the service, the associated port or protocol is not processed by anomaly detection.

Table 41-2 Destination Port or Protocol Map Dialog Box (continued)

Element	Description
Override Scanner Settings	Whether to override the scanner settings for this service or protocol. You must select this option to enable the remaining fields on the dialog box.
Scanner Threshold	The scanner threshold for this port or protocol. The range is 5 to 1000. The default is 200.
Threshold Histogram table	The histograms for this port or protocol. If you leave the table empty, the default histograms are used. You can have up to three rows, for low, medium, and high numbers of destination addresses, with different threshold levels (source addresses) for each. • To add a threshold, click the **Add Row** button and fill in the Histogram Dialog Box, page 41-13. The Add button is disabled if you already have three rows. • To edit a threshold, select it and click the **Edit Row** button. You cannot change the destination bucket to one that is already defined in the table. • To delete a threshold, select it and click the **Delete Row** button. Any buckets not included in the table use the default histogram for the bucket.

Histogram Dialog Box

Use the Histogram dialog box to create or modify entries in a histogram. The histograms you create or modify override the default histograms that anomaly detection generates. For detailed information about how these histograms are used, see:

- Understanding Anomaly Detection Thresholds and Histograms, page 41-9
- Configuring Anomaly Detection Thresholds and Histograms, page 41-11

Navigation Path

Do one of the following from the Anomaly Detection policy (see Configuring Anomaly Detection, page 41-6):

- On the TCP Protocol, UDP Protocol, or Other Protocol sub tabs on the Internal Zone, Illegal Zone, or External Zone tabs, select a row in the Threshold Histogram table and click the **Edit Row** button.
- On the Add or Modify Dest or Protocol Map dialog boxes, click the **Add Row** button, or select a row and click the **Edit Row** button. For information on opening the map dialog boxes, see Dest Port or Protocol Map Dialog Box, page 41-12.

Field Reference

Table 41-3 Histogram Dialog Box

Element	Description
Number of Destination IP Addresses	The histogram bucket you are defining. The buckets have a fixed number of destination addresses: Low (5 addresses); Medium (20); High (100). **Tip** A histogram can have a single entry for each destination bucket (low, medium, high). Thus, you cannot change this value to one that is already defined in the histogram you are editing.
Number of Source IP Addresses	The number of source addresses that are allowed to simultaneously scan the associated number of destination addresses. Enter the desired number. The range is 0 to 4096. If you are editing a histogram, the current value is shown.

Configuring Global Correlation

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

You can configure global correlation so that your sensors are aware of network devices with a reputation for malicious activity and can take action against them. Global correlation allows you to dynamically use information about malicious activity collected from networks around the globe to change the risk rating of events that have known bad devices as their source.

To configure global correlation, your sensor must be running IPS 7.0+ software. Global correlation is not available on Cisco IOS IPS devices.

This chapter contains the following topics:

- Understanding Global Correlation, page 42-1
- Configuring Global Correlation Inspection and Reputation, page 42-5
- Configuring Network Participation, page 42-7

Understanding Global Correlation

You can configure global correlation so that your sensors are aware of network devices with a reputation for malicious activity and can take action against them. Participating IPS devices in a centralized Cisco threat database, the SensorBase, receive and absorb global correlation updates. The reputation data contained in the global correlation updates is factored into the analysis of network traffic, which increases IPS efficacy, because traffic is denied or allowed based on the reputation of the source IP address. The participating IPS devices send data back to the Cisco SensorBase Network, which results in a feedback loop that keeps the updates current and global.

Tip The Botnet Traffic Filter feature of adaptive security appliances (ASA) is another dynamic feature you can deploy in your network to defend against malicious activity. Configuring global correlation on IPS devices, and Botnet Traffic Filtering on ASA firewalls, can be an effective combined security implementation. For more information about Botnet Traffic Filtering, see Chapter 19, “Managing Firewall Botnet Traffic Filter Rules”.

There are three main features of global correlation:

- Global Correlation Inspection—The IPS uses the global correlation reputation knowledge of attackers to influence alert handling and to deny actions when attackers with a bad score are seen on the sensor. For more information about reputation, see Understanding Reputation, page 42-2.
- Reputation Filtering—Applies automatic deny actions to packets from known malicious sites.
- Network Participation—The sensor sends alert and TCP fingerprint data to the SensorBase Network so that other users can share in the community knowledge. For more information, see Understanding Network Participation, page 42-3.

Global correlation has the following goals:

- Dealing intelligently with alerts thus improving efficacy.
- Improving protection against known malicious sites.
- Sharing telemetry data with the SensorBase Network to improve visibility of alerts and sensor actions on a global scale.
- Simplifying configuration settings.
- Automatic handling of the uploads and downloads of the information.

Tip You can use Report Manager to generate reports comparing the number of alerts generated by global correlation to those generated by traditional IPS inspection. For information on the Inspection/Global Correlation report, see Understanding General IPS Reports, page 70-19. For information on generating reports, see Opening and Generating Reports, page 70-20.

For information on how to configure global correlation, see the following topics:

- Global Correlation Requirements and Limitations, page 42-4
- Configuring Global Correlation Inspection and Reputation, page 42-5
- Configuring Network Participation, page 42-7

Understanding Reputation

Similar to human social interaction, reputation is an opinion toward a device on the Internet. Reputation indicates the probability that a particular attacker IP address will initiate malicious behavior based on its known past activity. Reputation enables the installed base of IPS sensors to collaborate using the existing network infrastructure and identify network devices that are likely to be malicious or infected.

By collecting data about devices and assigning reputation scores to them, the global correlation database provides important data that the IPS sensor can use to adjust the risk rating of an attack. Risk rating is the probability that a network event is malicious. Each signature has an associated risk rating. If you enable global correlation, the IPS sensor computes a score based on the reputation of an attacker and adds this score to the risk rating of the event. The updated risk rating is then used by your event action override and filter policies to help determine what actions to take for the event.

Thus, you might have an event that is initially configured to simply produce an alert. But, if the attacker has a bad reputation, the IPS might increase the risk rating to a number high enough that it triggers an event action override rule that adds the Deny Packet Inline action. Thus, for some source devices, the event simply produces an alert, but for others, the event drops the packet in addition to producing the alert.

Tip The Produce Alert action is added to an event whenever global correlation raises the risk rating of the event, or when global correlation adds the Deny Packet Inline or Deny Attacker Inline actions.

Because the global correlation database changes rapidly, the sensor must periodically download global correlation updates from the global correlation servers.

Using reputation scores to adjust the risk rating of an event improves the efficacy of the sensor by improving the following metrics:

- False positives as a percentage of actionable events.
- False negatives as a percentage of threats that do not result in actionable events.
- Actionable events as a percentage of all events.

Related Topics

- Understanding Global Correlation, page 42-1
- Understanding Network Participation, page 42-3
- Configuring Global Correlation Inspection and Reputation, page 42-5
- Configuring Network Participation, page 42-7

Understanding Network Participation

Network participation lets Cisco collect nearly real-time data from sensors around the world. Sensors installed at customer sites can send data to the SensorBase Network. These data feed in to the global correlation database to increase reputation fidelity. Communication between sensors and the SensorBase Network involves an HTTPS request and response over TCP/IP.

There are three modes for Network Participation:

- **Off**—The Network Participation server does not collect data, track statistics, or try to contact the Cisco SensorBase network.
- **Partial Participation**—The Network Participation server collects data, tracks statistics, and communicates with the SensorBase network. Data considered to be potentially sensitive is filtered out and never sent.

  Note Configuring the sensor for partial network participation limits a third party from extracting reconnaissance information about your internal network from the global correlation database.

- **Full Participation**—The Network Participation server collects data, tracks statistics, and communicates with the SensorBase network. All data collected is sent.

If you select partial or full participation, you are prompted to accept a participation agreement. You must accept the agreement to participate or you cannot change the participation mode.

The following table explains the data collected and the purpose of collecting it.

Table 42-1 Network Participation Data Sharing and Usage

Participation Level	Type of Data	Purpose
Partial	Protocol attributes (TCP maximum segment size and options string, for example).	Tracks potential threats and helps Cisco to understand threat exposure.
	Attack type (signature fired, including signature ID and version, risk rating, and reputation, for example).	Used to understand current attacks and attack severity.
	Connecting IP address and port.	Identifies attack source.
	Summary IPS performance (CPU utilization, memory usage, inline vs promiscuous, for example).	Tracks product efficacy.
Full	Victim IP address and port.	Detects threat behavioral pattern.

To configure network participation, the IPS device requires at least 100 MB of available memory, a network connection to the sensor, and a network connection to the Internet. For information on configuring network participation, see Configuring Network Participation, page 42-7.

Global Correlation Requirements and Limitations

The following list explains the requirements that you must meet to configure and successfully use global correlation on IPS devices. It also explains some limitations.

- **Valid license**—You must have a valid sensor license for global correlation features to function. You can still configure and display statistics for the global correlation features, but the global correlation databases are cleared and no updates are attempted. Once you install a valid license, the global correlation features are reactivated. For information on configuring licenses, see Updating IPS License Files, page 44-1.
- **Agree to Network Participation disclaimer**—If you decide to configure network participation, you must accept the disclaimer. For more information, see Understanding Network Participation, page 42-3 and Configuring Network Participation, page 42-7.
- **External connectivity for sensor and a DNS server or HTTP proxy**—Global correlation requires the sensor to connect to the Cisco SensorBase Network. Domain name resolution is also required for these features to function. You can either configure the sensor to connect through an HTTP proxy server that has a DNS client running on it, or you can assign an Internet routeable address to the management interface of the sensor and configure the sensor to use a DNS server. For more information, see Identifying DNS Servers, page 36-24 and Identifying an HTTP Proxy Server, page 36-24.
- **Sensor in inline mode**—The sensor must operate in inline mode so that the global correlation features can increase efficacy by being able to use the inline deny actions.
- **Sensor and IPS version that supports the global correlation features**—The sensor must run IPS 7.0+ software. You cannot configure global correlation on Cisco IOS IPS devices.
- **Sufficient available memory**—To configure network participation, the IPS device requires at least 100 MB of available memory.

- **Firewall access for port 80, 443 traffic**—Because global correlation updates occur through the sensor management interface, any firewall that lies between the sensor and the internet must allow traffic on ports 80 and 443. You can also use an HTTP proxy (see Identifying an HTTP Proxy Server, page 36-24).
- **Exposure to external traffic**—The global correlation database contains external IP addresses only, so if you position a sensor in an internal lab that has no interaction with outside networks, you might never receive global correlation information. The feature will have no effect.
- **Bypass mode might be triggered during global correlation updates**— As with signature updates, when the sensor applies a global correlation update, it might trigger bypass. Whether bypass is triggered depends on the traffic load of the sensor and the size of the signature or global correlation update. If bypass mode is turned off, an inline sensor stops passing traffic while the update is being applied.
- **No IPv6 address support**—Global correlation inspection and the reputation filtering deny features do not support IPv6 addresses. For global correlation inspection, the sensor does not receive or process reputation data for IPv6 addresses. The risk rating for IPv6 addresses is not modified for global correlation inspection. Similarly, network participation does not include event data for attacks from IPv6 addresses. And finally, IPv6 addresses do not appear in the deny list.

Related Topics

- Understanding Global Correlation, page 42-1
- Understanding Reputation, page 42-2
- Understanding Network Participation, page 42-3
- Configuring Global Correlation Inspection and Reputation, page 42-5
- Configuring Network Participation, page 42-7

Configuring Global Correlation Inspection and Reputation

Use the Inspection/Reputation policy to configure the sensor to use updates from the SensorBase Network to adjust the risk rating of events. The global correlation client on the sensor determines which updates are available and applicable to the sensor by communicating with the global correlation update server and a file server. The global correlation update server provides the server manifest document to the sensor, which identifies which updates are available and how to obtain them from a file server. The sensor downloads the update files from the file server using the information in the server manifest.

When you configure global correlation, updates are automatic and happen at regular intervals, approximately every five minutes by default, but this interval can be modified by the global correlation server. The sensor initially gets a full update and then applies an incremental update periodically.

If you turn on global correlation, you can choose how aggressively you want the deny actions to be enforced against malicious hosts. You can then enable reputation filtering to deny access to known malicious hosts. If you only want a report of what could have happened, you can enable Test Global Correlation. This puts the sensor in Audit mode, and actions the sensor would have performed are generated in the events.

Tip When you view IPS events in Event Viewer, there are several columns specific to global correlation that you can add to the event table; these columns are not shown by default, so you must add them to your view. To monitor global correlation in general, use the IPS device manager (IDM) and look at the Sensor Health gadget. Use either the full IDM or open a read-only copy from Security Manager by right-clicking the device in Device view and selecting **Device Manager**.

Before You Begin

- You must also configure a DNS server or HTTP proxy for global correlation to function. For details, see Identifying DNS Servers, page 36-24 or Identifying an HTTP Proxy Server, page 36-24.
- There are several configuration requirements and limitations that you should be aware of before configuring global configuration. For details, see Global Correlation Requirements and Limitations, page 42-4.

Related Topics

- Understanding Global Correlation, page 42-1
- Understanding Reputation, page 42-2
- Configuring Network Participation, page 42-7

Step 1 Do one of the following to open the Inspection/Reputation policy:

- (Device view) Select **IPS > Global Correlation > Inspection/Reputation** from the Policy selector.
- (Policy view) Select **IPS > Global Correlation > Inspection/Reputation** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Configure the following settings:

- **Global Correlation Inspection**—Whether to enable global correlation inspection. When turned on, the sensor uses updates from the SensorBase Network to adjust the risk rating. Deselect this option to disable inspection.
- **Global Correlation Influence**—How aggressively the sensor uses global correlation information to initiate deny actions. Select one of the following:
  - **Permissive**—Has the least aggressive effect on deny actions.
  - **Standard**—(The default.) Has a moderately aggressive effect on deny actions.
  - **Aggressive**—Has a very aggressive effect on deny actions.
- **Reputation Filtering**—Select whether you want reputation filtering **on** or **off**. When turned on, the sensor denies access to malicious hosts that are listed in the global correlation database.
- **Test Global Correlation**—Whether to place global correlation in audit mode. In audit mode, reputation filtering does not deny access to known malicious hosts; only a report of what could have happened is generated.

  Audit mode allows you to test the global correlation features without actually denying any hosts. If you decide the effects are desirable, you can deselect this option to activate reputation filtering.

Configuring Network Participation

Use the Network Participation policy to configure the sensor to send data to the SensorBase Network. You can configure the sensor to fully participate and send all data to the SensorBase Network, or you can configure the sensor to collect the data but to omit potentially sensitive data, such as the destination IP address of trigger packets. For detailed information about network participation and the data that is collected, see Understanding Network Participation, page 42-3.

Related Topics

- Understanding Global Correlation, page 42-1
- Understanding Reputation, page 42-2
- Global Correlation Requirements and Limitations, page 42-4
- Configuring Global Correlation Inspection and Reputation, page 42-5

Step 1 Do one of the following to open the Network Participation policy:

- (Device view) Select **IPS > Global Correlation > Network Participation** from the Policy selector.
- (Policy view) Select **IPS > Global Correlation > Network Participation** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Select the level of participation from the **Network Participation** list:

- **Off**—No data is contributed to the SensorBase network.
- **Partial**—Data is contributed to the SensorBase network but potentially sensitive information is withheld.

  Note Configuring the sensor for partial network participation limits a third party from extracting reconnaissance information about your internal network from the Global Correlation database.

- **Full**—All data is contributed to the SensorBase network.

Step 3 If you select Full or Partial, when you click Save, the Network Participation Disclaimer dialog box opens, prompting you to read and accept a disclaimer. Carefully read the disclaimer. Click **Agree** if you agree to it.

If you click **Disagree**, you cannot enable network participation. Change the setting to Off and save the policy.

Configuring Attack Response Controller for Blocking and Rate Limiting

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any enhancements as IPS is now End of Life. For more information, see EOL notice.

You can configure an IPS device to implement blocks or rate limits to control attacks. Blocking and rate limiting are primarily of use when operating in promiscuous mode. When operating in inline mode, it is much more efficient to have the IPS drop traffic itself. Blocking and rate limiting are actions that other devices implement at the request of the IPS; thus, configuring blocking and rate limiting is a more complex configuration than simple inline denies.

To configure blocking or rate limiting, you must identify the network device that performs the blocking. A network device that performs blocking is called a blocking device. Many network devices can be used to support blocking: Cisco IOS routers and Catalyst 6500 switches, Cisco security appliances (ASA, PIX, and FWSM), and Catalyst 6500/7600 devices running the Catalyst operating system. You can also configure another IPS device to act as a main blocking sensor.

Note IPS blocking and rate limiting works only for IPS appliances and service modules. You cannot configure it for Cisco IOS IPS.

This chapter contains the following topics:

- Understanding IPS Blocking, page 43-1
- Configuring IPS Blocking and Rate Limiting, page 43-7
- Blocking Page, page 43-8

Understanding IPS Blocking

The Attack Response Controller (ARC) component of the IPS is responsible for managing network devices in response to suspicious events by blocking access from attacking hosts and networks. ARC blocks the IP address on the devices it is managing. It sends the same block to all the devices it is managing, including any other main blocking sensors. ARC monitors the time for the block and removes the block after the time has expired.

Note ARC is formerly known as Network Access Controller. Although the name has been changed, the IPS documentation and configuration interfaces contain references to Network Access Controller, nac, and network-access.

ARC completes the action response for a new block in no more than 7 seconds. In most cases, it completes the action response in less time. To meet this performance goal, you should not configure the sensor to perform blocks at too high a rate or to manage too many blocking devices and interfaces. We recommend that the maximum number of blocks not exceed 250 and the maximum number of blocking items not exceed 10. To calculate the maximum number of blocking items, a security appliance counts as one blocking item per blocking context. A router counts as one blocking item per blocking interface/direction. A switch running Catalyst software counts as one blocking item per blocking VLAN. If the recommended limits are exceeded, ARC might not apply blocks in a timely manner or might not be able to apply blocks at all.

For security appliances configured in multiple-context mode, Cisco IPS does not include VLAN information in the block request. Therefore you must make sure the IP addresses being blocked are correct for each security appliance. For example, the sensor is monitoring packets on a security appliance customer context that is configured for VLAN A, but is blocking on a different security appliance customer context that is configured for VLAN B. Addresses that trigger blocks on VLAN A might refer to a different host on VLAN B.

Note Blocking is not supported on the FWSM on the admin context in multiple-context mode.

There are three types of blocks:

- Host block—Blocks all traffic from a given IP address.

  To configure the IPS to initiate automatic host blocks when a signature is triggered, add the **Request Block Host** event action to a signature, or add it to events based on risk rating using the event action override policy. See Configuring Event Action Overrides, page 40-13 and Configuring Signatures, page 39-4.

- Connection block—Blocks traffic from a given source IP address to a given destination IP address and destination port. Multiple connection blocks from the same source IP address to either a different destination IP address or destination port automatically switch the block from a connection block to a host block.

  To configure the IPS to initiate automatic connection blocks when a signature is triggered, add the **Request Block Connection** event action to a signature, or add it to events based on risk rating using the event action override policy.

- Network block—Blocks all traffic from a given network.

  You can initiate host and connection blocks manually or automatically when a signature is triggered. You can only initiate network blocks manually. You cannot initiate network blocks from within Security Manager; use the IPS Device Manager instead.

Tip Connection blocks and network blocks are not supported on security appliances (firewalls). Security appliances only support host blocks with additional connection information.

Note Do not confuse blocking with the ability of the sensor to drop packets. The sensor can drop packets when the following actions are configured for a sensor in inline mode: deny packet inline, deny connection inline, and deny attacker inline.

On Cisco IOS Software devices (routers and Catalyst 6500 series switches), ARC creates blocks by applying ACLs; on Catalyst 6500/7600 devices that run the Catalyst operating system, ARC creates blocks by applying VACLs. ACLs and VACLs permit or deny passage of data packets through interface directions or VLANs. Each ACL or VACL contains permit and deny conditions that apply to IP addresses. The security appliances use the **shun** command instead of ACLs.

Tip For a list of the specific devices and operating system versions that you can configure as blocking devices, see the supported device information in the chapter "Configuring Attack Response Controller for Blocking and Rate Limiting" in the *Installing and Using Cisco Intrusion Prevention System Device Manager* publication for your IPS software version. These publications are available on Cisco.com at http://www.cisco.com/en/US/products/hw/vpndevc/ps4077/products_installation_and_configuration_guides_list.html.

The following topics explain more about IPS blocking:

- Strategies for Applying Blocks, page 43-3
- Understanding Rate Limiting, page 43-4
- Understanding Router and Switch Blocking Devices, page 43-4
- Understanding the Master Blocking Sensor, page 43-6
- Configuring IPS Blocking and Rate Limiting, page 43-7
- Blocking Page, page 43-8

Strategies for Applying Blocks

Blocking is performed only when an event occurs and the event includes the Request Block Connection or Request Block Host event actions. These event actions are not typically needed when you operate the IPS in inline mode, where you use Deny actions to drop undesired traffic.

The following are situations in which you might want to implement blocking actions:

- Promiscuous mode—When running in promiscuous mode, the IPS cannot implement Deny actions. Thus, if you want to prevent traffic from a host, you must implement blocking.
- Inline mode—In inline mode, you can implement Deny actions to immediately drop undesired traffic. However, you might want to add blocking actions to protect other segments of your network.

  For example, suppose that your network consists of five subnets, A, B, C, D, and E, and that each of these segments has an inline IPS device monitoring it. If the IPS for subnet A identifies an attack, the IPS can use Deny actions to protect subnet A, but also use Request Block actions to configure the firewalls that protect B, C, D, and E to shun the attacker before the attack can target those other subnets. In this example, you would want to designate a single IPS as the main blocking sensor and have the other four IPS sensors perform blocking through the main blocking sensor.

Use the following techniques to add the request block actions to an event:

- Event Action Override policy—Configure an event action override rule to add the action to all events based on the event's risk rating. This is a simple approach. You could add the request block action for the same risk ratings used for adding Deny actions. For more information, see Configuring Event Action Overrides, page 40-13.
- Signatures policy—You can add the request block actions to individual signatures. This requires editing each signature to add the action. This can be a time-consuming approach, but it allows you to configure blocking for just the types of events that concern you most. For more information, see Configuring Signatures, page 39-4.

Related Topics

- Understanding IPS Blocking, page 43-1
- Understanding the Master Blocking Sensor, page 43-6
- Understanding Interface Modes, page 37-2
- Configuring IPS Blocking and Rate Limiting, page 43-7
- Blocking Page, page 43-8

Understanding Rate Limiting

Attack Response Controller (ARC) is responsible for rate limiting traffic in protected networks. Rate limiting lets sensors restrict the rate of specified traffic classes on network devices. Rate limit responses are supported for the Host Flood and Net Flood engines, and the TCP half-open SYN signature. ARC can configure rate limits on network devices running Cisco IOS 12.3 or later. Main blocking sensors can also forward rate limit requests to blocking forwarding sensors.

To add a rate limit to a signature, you must add the Request Rate Limit action. You can then edit the signature parameters to set the percentage for these signatures in the Event Actions Settings folder.

Tip You can also manually implement rate limits, but you cannot do so using Security Manager; use the IPS Device Manager instead.

On the blocking device, you must not apply a service policy to an interface/direction that is configured for rate limiting. If you do so, the rate limit action will fail. Before configuring rate limits, confirm that there is no service policy on the interface/direction, and remove it if one exists. ARC does not remove the existing rate limit unless it is one that ARC had previously added.

Rate limits use ACLs, but not in the same way as blocks. Rate limits use ACLs and class-map entries to identify traffic, and policy-map and service-policy entries to police the traffic.

Understanding Router and Switch Blocking Devices

You can use routers or Catalyst 6500/7600 devices running Cisco IOS Software, or Catalyst 6500/7600 devices running the Catalyst operating system, to implement IPS blocking in your network. When you use routers or switches, Attack Response Controller (ARC) configures extended ACLs (on IOS devices) or VLAN ACLs (on Catalyst OS devices) to implement the blocks. These ACLs and VACLs are created and managed in the same way.

Rate limits also use ACLs, but not in the same way as blocks. Rate limits use ACLs and class-map entries to identify traffic, and policy-map and service-policy entries to police the traffic.

Tip IPS considers Catalyst 6500/7600 devices that run Cisco IOS Software to be equivalent to routers. When you add these devices as blocking devices, add them as routers.

When you configure a router interface or switch VLAN as a blocking interface, you can optionally specify the names of pre- and post-ACLs or VACLs. Although specifying ACL or VACL names is optional, if you have configured ACLs or VACLs on the interface or VLAN, you must identify them to the IPS or ARC will remove them from your device configuration.

The pre- and post-ACL/VACL have the following uses:

- The Pre-Block ACL/VACL is mainly used for permitting what you do not want the sensor to ever block. When a packet is checked against the ACL/VACL, the first line that gets matched determines the action. If the first line matched is a permit line from the Pre-Block ACL/VACL, the packet is permitted even though there may be a deny line (from an automatic block) listed later in the ACL/VACL. The Pre-Block ACL/VACL can override the deny lines resulting from the blocks.
- The Post-Block ACL/VACL is best used for additional blocking or permitting that you want to occur on the same interface or direction. If you have an existing ACL on the interface or direction that the sensor will manage, that existing ACL can be used as a Post-Block ACL/VACL. If you do not have a Post-Block ACL/VACL, the sensor inserts permit ip any any at the end of the new ACL/VACL.

  If you are managing the IOS Software blocking device in Security Manager, you can identify the ACL name by selecting the blocking device, then selecting **Tools > Preview Config**. Look for the **ip access-group** command in the interface configuration, and check the direction. For example, the following lines show that there is an ACL named CSM_FW_ACL_GigabitEthernet0/1 in the In direction attached to the GigabitEthernet0/1 interface.

  ```
  interface GigabitEthernet0/1
    ip access-group CSM_FW_ACL_GigabitEthernet0/1 in
  ```

  In this example, if you configure GigabitEthernet0/1 in the In direction as a blocking interface, ensure that you specify CSM_FW_ACL_GigabitEthernet0/1 as a pre- or post-ACL. In most cases, you should specify the ACL as the post-ACL, so that the relatively short IPS blocking ACL first filters out undesirable traffic before the blocking device implements your other access rules.

  Because Security Manager does not manage Catalyst OS devices, you must examine a Catalyst OS device configuration outside of Security Manager to determine VACL names. Keep in mind that a Catalyst 6500/7600 device that runs IOS Software can also have VACLs, but the IPS does not do VLAN blocking on Catalyst 6500/7600 VLANs when the device is running IOS Software.

When the sensor starts up, it reads the contents of the two ACL/VACLs. It creates a third ACL/VACL with the following entries in this order, and this combined ACL/VACL is applied to the interface or VLAN:

1. A **permit** line with the sensor IP address or, if specified, the NAT address of the sensor.

   If you select the Allow Sensor IP address to be Blocked option on the General tab of the Blocking policy, this permit entry is not added. For more information, see General Tab, IPS Blocking Policy, page 43-10.
2. Pre-Block ACL/VACL, if specified.
3. Any active blocks generated by the IPS (deny statements).
4. The Post-Block ACL/VACL, if specified.

   If you do not specify a Post-Block ACL/VACL, a **permit ip any any** entry is added to allow all unfiltered traffic. Note that this negates the normal implicit deny any that ends interface ACLs.

   When using Catalyst OS, IDSM-2 inserts **permit ip any any capture** at the end of the new VACL.

If ARC is managing a device and you need to configure the ACL/VACLs on that device, you should disable blocking first. You want to avoid a situation in which both you and ARC could be making a change at the same time on the same device. This could cause the device or ARC to fail. If you need to modify the Pre-Block or Post-Block ACL/VACL, do the following:

1. Disable blocking on the sensor.

   Because you are making a temporary change, you can disable and then reenable blocking by using the IPS Device Manager (IDM) on the device. Alternatively, you can deselect the Enable Blocking option on the General tab of the Blocking policy in Security Manager, then deploy the configuration to the IPS sensor. To reenable blocking, select the Enable Blocking option again and deploy the configuration to the IPS sensor.
2. Make the changes to the configuration of the device. For example, if you manage the blocking device in Security Manager, deploy the updated configuration and wait for the device to reload.
3. Reenable blocking on the sensor.

Understanding the Master Blocking Sensor

Multiple sensors (blocking forwarding sensors) can forward blocking requests to a specified main blocking sensor, which controls one or more devices. The main blocking sensor is the ARC running on a sensor that controls blocking on one or more devices on behalf of one or more other sensors. When a signature fires that has blocking or rate limit requests configured as event actions, the sensor forwards the block or rate limit request to the main blocking sensor, which then performs the block or rate limit.

When you add a main blocking sensor, you reduce the number of blocking devices per sensor. For example, if you want to block on 10 firewalls and 10 routers with one blocking interface/direction each, you can assign 10 to the sensor and assign the other 10 to a main blocking sensor.

You configure main blocking sensors on the Master Blocking Sensors tab of the Blocking policy, as described in Blocking Page, page 43-8.

When configuring the main blocking sensors, keep the following tips in mind:

- Two sensors cannot control blocking or rate limiting on the same device. If this situation is needed, configure one sensor as the main blocking sensor to manage the devices and the other sensors can forward their requests to the main blocking sensor.
- On the blocking forwarding sensor, identify which remote host serves as the main blocking sensor; on the main blocking sensor you must add the blocking forwarding sensors to its access list using the Allowed Hosts policy. See Identifying Allowed Hosts, page 36-7.
- If the main blocking sensor requires TLS for web connections, you must configure the ARC of the blocking forwarding sensor to accept the X.509 certificate of the main blocking sensor remote host. Sensors by default have TLS enabled, but you can change this option. For more information, see Master Blocking Sensor Dialog Box, page 43-13.
- Typically the main blocking sensor is configured to manage the network devices. Blocking forwarding sensors are not normally configured to manage other network devices, although doing so is permissible.
- Only one sensor should control all blocking interfaces on a device.

Configuring IPS Blocking and Rate Limiting

If you use the Request Block Host, Request Block Connection, or Request Rate Limit actions on any signatures, or add them to events using the event action override policy, you must configure blocking devices. If you do not use these actions, there is no need to configure blocking devices.

Before you configure blocking, read the following topics:

- Understanding IPS Blocking, page 43-1
- Strategies for Applying Blocks, page 43-3
- Understanding Rate Limiting, page 43-4
- Understanding Router and Switch Blocking Devices, page 43-4
- Understanding the Master Blocking Sensor, page 43-6

Step 1 Do one of the following:

- (Device view) Select **Platform > Security > Blocking** from the Policy selector.
- (Policy view) Select **IPS > Platform > Security > Blocking**, then select an existing policy or create a new one.

For an overview of the blocking policy, see Blocking Page, page 43-8.

Step 2 On the General tab, change any settings where you want non-default values. However, the default values are appropriate for most networks. For detailed information about the settings, see General Tab, IPS Blocking Policy, page 43-10.

Step 3 Click the **User Profiles** tab and create the user profiles that are required to log into the blocking devices.

- To add a profile, click the **Add Row** button and fill in the Add User Profile dialog box (see User Profile Dialog Box, page 43-12).
- To edit a profile, select it and click the **Edit Row** button.
- To delete a profile, select it and click the **Delete Row** button. Before you delete a profile, ensure that it is not currently being used by a blocking device.

Step 4 If you need to use a main blocking sensor, as described in Understanding the Master Blocking Sensor, page 43-6, click the **Master Blocking Sensors** tab and do the following:

- To add a main blocking sensor, click the **Add Row** button and fill in the Add Master Blocking Sensor dialog box (see Master Blocking Sensor Dialog Box, page 43-13).
- To edit a main blocking sensor, select it and click the **Edit Row** button.
- To delete a main blocking sensor, select it and click the **Delete Row** button.

Step 5 Identify the blocking devices (unless you will use main blocking sensors only). You must add the devices to the correct tab:

- **Routers** tab—For all Cisco IOS Software devices, including Catalyst 6500 switches that are running IOS Software.
- **Firewalls** tab—For ASA, PIX, and FWSM.
- **Catalyst 6K** tab—For Catalyst 6500/7600 devices that are running the Catalyst operating system.

On each tab, the configuration steps are the same:

- To add a device, click the **Add Row** button and fill in the Add Router, Firewall, or Cat6K Device dialog box (see Router, Firewall, Cat6K Device Dialog Box, page 43-14).

- To edit a device, select it and click the **Edit Row** button.
- To delete a device, select it and click the **Delete Row** button.

Step 6 Click the **Never Block Hosts and Networks** tab and identify the hosts and networks that should never be blocked. These lists affect blocking actions, but they do not affect limiting actions. Identify your trusted networks and hosts:

- To add a host or network, click the **Add Row** button beneath the appropriate table and fill in the Add Never Block Host or Network dialog box (see Never Block Host or Network Dialog Boxes, page 43-17).
- To edit a host or network, select it and click the **Edit Row** button.
- To delete a host or network, select it and click the **Delete Row** button.

Blocking Page

Use the Blocking page to configure IPS sensor blocking properties. Configure the blocking policy only if you use the Request Block Connection, Request Block Host, or Request Rate Limit event actions in your signatures or event actions policies. Blocking hosts are used only for events to which these actions are assigned.

Tip The list of hosts and networks to never block applies only to the Request Block Connection and Request Block Host event actions. The list does not affect rate limiting, nor does it affect any of the Deny actions such as Deny Packet Inline. To exempt hosts and networks from Deny or rate limiting actions, use event action filter rules, specify the hosts and networks as Attackers, and remove the actions from events. For more information, see Configuring Event Action Filters, page 40-4.

Navigation Path

- (Device view) Select **Platform > Security > Blocking** from the Policy selector.
- (Policy view) Select **IPS > Platform > Security > Blocking**, then select an existing policy or create a new one.

Related Topic

- Configuring IPS Blocking and Rate Limiting, page 43-7
- Understanding IPS Blocking, page 43-1
- Strategies for Applying Blocks, page 43-3
- Understanding Rate Limiting, page 43-4
- Understanding Router and Switch Blocking Devices, page 43-4
- Understanding the Master Blocking Sensor, page 43-6
- Understanding IPS Event Actions, page 40-2

Field Reference

Table 43-1 IPS Blocking Policy

Element	Description
General tab	The basic settings required to enable blocking and rate limiting. For information about the options on the General tab, see General Tab, IPS Blocking Policy, page 43-10.
User Profiles tab	The connection credential information profiles for logging into the blocking devices. Before you define a blocking device, create the user profile required to log into the device. The table shows the profile name, username, and the passwords, which are masked with a fixed number of asterisks. • To add a profile, click the **Add Row** button and fill in the Add User Profile dialog box (see User Profile Dialog Box, page 43-12). • To edit a profile, select it and click the **Edit Row** button. • To delete a profile, select it and click the **Delete Row** button. Before you delete a profile, ensure that it is not currently being used by a blocking device.
Master Blocking Sensors tab	The main blocking IPS sensors (see Understanding the Master Blocking Sensor, page 43-6). A main blocking sensor manages blocks for other IPS devices. The table shows the IP address (or network/host object) of the main blocking sensor, the username and password for logging into it, the port used for connections, and whether TLS is used for login. • To add a main blocking sensor, click the **Add Row** button and fill in the Add Master Blocking Sensor dialog box (see Master Blocking Sensor Dialog Box, page 43-13). • To edit a main blocking sensor, select it and click the **Edit Row** button. • To delete a main blocking sensor, select it and click the **Delete Row** button.
Router tab	The IOS routers and Catalyst 6500/7600 devices (that are running IOS Software) to be used as blocking or rate limiting devices. The table shows the IP address (or network/host object) of the device, the communication method used to log into it, the NAT address of the sensor (0.0.0.0 if NAT is not used), the name of the profile that is used for logging into the device, and the device's response capabilities (blocking, rate limiting, or both). • To add a router, click the **Add Row** button and fill in the Add Router Device dialog box (see Router, Firewall, Cat6K Device Dialog Box, page 43-14). • To edit a router, select it and click the **Edit Row** button. • To delete a router, select it and click the **Delete Row** button.

Table 43-1 IPS Blocking Policy (continued)

Element	Description
Firewall tab	The ASA, PIX, and FWSM devices to be used as blocking devices. The table shows the IP address (or network/host object) of the device, the communication method used to log into it, the NAT address of the sensor (0.0.0.0 if NAT is not used), and the name of the profile that is used for logging into the device. • To add a firewall, click the **Add Row** button and fill in the Add Firewall Device dialog box (see Router, Firewall, Cat6K Device Dialog Box, page 43-14). • To edit a firewall, select it and click the **Edit Row** button. • To delete a firewall, select it and click the **Delete Row** button.
Catalyst 6K tab	The Catalyst 6500/7600 devices that are using Catalyst software to be used as blocking devices. The table shows the IP address (or network/host object) of the device, the communication method used to log into it, the NAT address of the sensor (0.0.0.0 if NAT is not used), and the name of the profile that is used for logging into the device. **Tip** Do not use this tab for Catalyst 6500/7600 devices that run Cisco IOS Software. Instead, use the Router tab. • To add a Catalyst OS device, click the **Add Row** button and fill in the Add Cat6K Device dialog box (see Router, Firewall, Cat6K Device Dialog Box, page 43-14). • To edit a Catalyst OS device, select it and click the **Edit Row** button. • To delete a Catalyst OS device, select it and click the **Delete Row** button.
Never Block Hosts and Networks	The hosts and networks that should never be blocked. Hosts and networks are shown in separate tables. The tables show the IP address or network/host object for the host or network. These lists do not affect rate limiting actions, nor do they apply to Deny actions. • To add a host or network, click the **Add Row** button beneath the appropriate table and fill in the Add Never Block Host or Network dialog box (see Never Block Host or Network Dialog Boxes, page 43-17). • To edit a host or network, select it and click the **Edit Row** button. • To delete a host or network, select it and click the **Delete Row** button.

General Tab, IPS Blocking Policy

Use the General tab of the Blocking policy to configure the basic settings required to enable blocking and rate limiting.

Navigation Path

- (Device view) Select **Platform > Security > Blocking** from the Policy selector. If necessary, select the **General** tab.
- (Policy view) Select **IPS > Platform > Security > Blocking**, then select an existing policy or create a new one. If necessary, select the **General** tab.

Related Topic

- Understanding IPS Blocking, page 43-1
- Configuring IPS Blocking and Rate Limiting, page 43-7
- Blocking Page, page 43-8

Field Reference

Table 43-2 General Tab, IPS Blocking Policy

Element	Description
Log All Block Events and Errors	Whether to log events that follow blocks from start to finish and any error messages that occur. When a block is added to or removed from a device, an event is logged. You may not want all these events and errors to be logged. Disabling this option suppresses new events and errors. The default is enabled. **Note** Log all block events and errors also applies to rate limiting.
Enable NVRAM Write	Whether to have the router write to non-volatile RAM (NVRAM) when Attack Response Controller (ARC) first connects. If enabled, NVRAM is written each time the ACLs are updated. The default is disabled. Enabling NVRAM writing ensures that all changes for blocking and rate limiting are written to NVRAM. If the router is rebooted, the correct blocks and rate limits will still be active. If NVRAM writing is disabled, a short time without blocking or rate limiting occurs after a router reboot. Not enabling NVRAM writing increases the life of the NVRAM and decreases the time for new blocks and rate limits to be configured.
Enable ACL Logging	Whether to have ARC append the log parameter to block entries in the access control list (ACL) or VLAN ACL (VACL). This causes the device to generate syslog events when packets are filtered. This option applies to routers and switches only. The default is disabled.
Allow Sensor IP address to be Blocked	Whether the sensor IP address can be blocked. The default is disabled. **Tip** If you allow the sensor address to be blocked, the IPS does not add an explicit permit entry to the interface ACL to allow the IPS address. You must ensure that the IPS address is permitted by the device ACL or the IPS cannot implement blocking on the device.
Enable Blocking	Whether to enable the blocking and rate limiting of hosts. The default is enabled. **Note** When you enable blocking, you also enable rate limiting. When you disable blocking, you also disable rate limiting. This means that ARC cannot add new or remove existing blocks or rate limits.

Table 43-2 General Tab, IPS Blocking Policy (continued)

Element	Description
Max Blocks	The maximum number of entries to block. The range is 1 to 65535. The default is 250.
Max Interfaces	The maximum number of interfaces for performing blocks. For example, a PIX 500 series security appliance counts as one interface. A router with one interface counts as one, but a router with two interfaces counts as two. The maximum number of interfaces is 250 per device. The default is 250. You use Max Interfaces to set an upper limit on the number of devices and interfaces that ARC can manage. The total number of blocking devices (not including main blocking sensors) cannot exceed this value. The total number of blocking items also cannot exceed this value, where a blocking item is one security appliance context, one router blocking interface/direction, or one Catalyst Software switch blocking VLAN. **Note** In addition, the following maximum limits are fixed and you cannot change them: 100 interfaces per device, 250 security appliances, 250 routers, 250 Catalyst Software switches, and 100 main blocking sensors.
Max Rate Limits	The maximum number of rate limit entries. The maximum rate limit must be equal to or less than the maximum blocking entries. The range is 1 to 32767. The default value is 250.

User Profile Dialog Box

Use the Add or Modify User Profile dialog box to add or modify a user profile for an IPS blocking device. The profile defines a username and passwords that the IPS device can use to log into and configure the router, switch, or firewall that will implement IPS blocking.

Although you can save a profile that has a profile name only, the requirements for username, password, and enable password are determined by the device. You must specify the items required by the device to enter configuration mode, or the IPS cannot configure blocking on the device.

Navigation Path

From the IPS Blocking policy, select the User Profiles tab and click the **Add Row** button or select an existing sensor and click the **Edit Row** button. For information on opening the Blocking policy, see Blocking Page, page 43-8.

Field Reference

Table 43-3 User Profile Dialog Box

Element	Description
Profile Name	The name of the profile, up to 64 alphanumeric characters.
Username	The username to use when logging into the blocking device.
Password	The login password for the username, if required.

Table 43-3 User Profile Dialog Box (continued)

Element	Description
Enable Password	The enable password for entering Privileged EXEC Mode (enable mode), if required.

Master Blocking Sensor Dialog Box

Use the Add or Modify Master Blocking Sensor dialog box to configure a main blocking sensor. For more information about main blocking sensors, see Understanding the Master Blocking Sensor, page 43-6.

Navigation Path

From the IPS Blocking policy, select the Master Blocking Sensors tab and click the **Add Row** button or select an existing sensor and click the **Edit Row** button. For information on opening the Blocking policy, see Blocking Page, page 43-8.

Field Reference

Table 43-4 Master Blocking Sensor Dialog Box

Element	Description
IP Address	The IP address of the main blocking sensor. Enter the IP address or the name of a network/host policy object that contains a single host address, or click **Select** to select an object from a list or to create a new one.
Username	The username to use to log in to the main blocking sensor. The user account must be an active account configured on the main blocking sensor.
Password	The login password for the username.
Port	The port on which to connect on the main blocking sensor. The default is 443.
TLS	Whether to use TLS. If you select the TLS option, you must configure the ARC of the blocking forwarding sensor to accept the TLS/SSL X.509 certificate of the main blocking sensor remote host. (The blocking forwarding sensor is any device to which you are assigning this blocking policy.) The easiest way to configure the blocking forwarding sensor to accept the X.509 certificate is to use the IPS Device Manager (IDM) to log into the sensor, choose **Configuration > Sensor Management > Certificates > Trusted Hosts > Add Trusted Host**, and add the main blocking sensor as a trusted host. Alternatively, you can log into the sensor CLI, enter configuration mode, and use the **tls trusted-host ip-address** command.

Router, Firewall, Cat6K Device Dialog Box

Use the Add or Modify Router, Firewall, or Cat6K Device dialog box to configure a device as a blocking device for an IPS sensor. The name of the dialog box indicates the type of device you are adding:

- Router—IOS Software routers and Catalyst 6500/7600 devices. These devices can do rate limiting as well as blocking. See Understanding Router and Switch Blocking Devices, page 43-4.
- Firewall—ASA and PIX appliances.
- Cat6K—Catalyst 6500/7600 devices that are running Catalyst OS software.

Tip If the Catalyst 6500/7600 runs Cisco IOS Software, add the device as a router on the Router tab. Do not add the device to the Cat6K tab.

Navigation Path

From the IPS Blocking policy, select the Router, Firewall, or Catalyst 6K tab and click the **Add Row** button or select an existing row and click the **Edit Row** button. For information on opening the Blocking policy, see Blocking Page, page 43-8.

Field Reference

Table 43-5 Router, Firewall, Cat6K Device Dialog Boxes

Element	Description
IP Address	The IP address of the device. Enter the IP address or the name of a network/host policy object that contains a single host address, or click **Select** to select an object from a list or to create a new one.
Communication Type	The communication mechanism used to log in to the blocking device (SSH 3DES, SSH DES, Telnet). The default is SSH 3DES. If you choose SSH 3DES or SSH DES, you must add the device to the known hosts list. The easiest way to add the device to the known hosts list is to use the IPS Device Manager (IDM) to log into the sensor, choose **Configuration > Sensor Management > SSH > Known Host Keys > Add Known Host Key**, and add the device address. Alternatively, you can log into the sensor CLI, enter configuration mode, and use the **ssh host-key** command.
NAT Address	The NAT address of the sensor, if any is used between the sensor and the blocking device. Enter the NAT address or the name of a network/host policy object that contains a single host address, or click **Select** to select an object from a list or to create a new one. Leave the default 0.0.0.0 if NAT is not used.
Profile Name	The login profile used to log in to the blocking device. You must create this profile on the User Profiles tab of the blocking policy or the IPS cannot successfully use this blocking device.

Table 43-5 Router, Firewall, Cat6K Device Dialog Boxes (continued)

Element	Description
Interfaces and directions where blocks will be applied (table) (Routers only.)	The interfaces on the device that should be used for blocking or rate limiting. The table shows the interface name, direction, and the names of existing ACLs that the IPS device should incorporate into the blocking ACL. If the interface already has an ACL configured for the specified direction, you must specify that ACL name as a pre- or post-ACL or the IPS removes the ACL. These ACLs are used for blocking only, not for rate limiting. • To add an interface, click the **Add Row** button and fill in the Add Router Block Interface dialog box (see Router Block Interface Dialog Box, page 43-15). • To edit an interface, select it and click the **Edit Row** button. • To delete an interface, select it and click the **Delete Row** button.
Response Capabilities (Routers only.)	The actions that this router can implement. Use Ctrl+click to select multiple actions (highlighted actions are selected). Options are: • **Block**—The router can implement blocks in response to Request Block Connection and Request Block Host actions. • **Rate Limit**—The router can implement rate limits in response to Request Rate Limit actions.
VLANs where blocks will be applied (table) (Catalyst 6500/7600 devices running the Catalyst operating system only.)	The VLANs on the device that should be used for blocking. The table shows the VLAN name and the names of existing VLAN ACLs (VACL) that the IPS device should incorporate into the blocking VACL. If the VLAN already has a VACL configured, you must specify that VACL name as a pre- or post-VACL or the IPS removes the VACL. • To add a VLAN, click the **Add Row** button and fill in the Add Cat6K Block VLAN dialog box (see Cat6k Block VLAN Dialog Box, page 43-16). • To edit a VLAN, select it and click the **Edit Row** button. • To delete a VLAN, select it and click the **Delete Row** button.

Router Block Interface Dialog Box

Use the Add or Modify Router Block Interface dialog box to configure a blocking interface on a router or IOS Software Catalyst 6500/7600 device that is configured as an IPS blocking device. The IPS sensor uses the interface for blocking actions.

Navigation Path

From the Add or Modify Router Device dialog box, click the **Add Row** button beneath the interfaces table, or select a row in the table and click the **Edit Row** button. For information on opening the Router Device dialog box, see Router, Firewall, Cat6K Device Dialog Box, page 43-14.

Field Reference

Table 43-6 Router Block Interface Dialog Box

Element	Description
Interface Name	The name of the interface on the router that the IPS should use for blocking. Enter the name exactly as it is configured on the router (for example, GigabitEthernet0/1).
Direction	The direction to apply the blocking ACL, In or Out.

Pre ACL Name
Post ACL Name
	The ACLs to combine with the blocking entries that the IPS creates to implement blocking actions. The Pre ACL is added before the blocking ACL, and the Post ACL is added after the blocking ACL. For more information, see Understanding Router and Switch Blocking Devices, page 43-4.
	Tip If you have configured an ACL on the interface in the specified direction, you must specify the name of the ACL in the Pre or Post ACL Name field or the ACL will be removed from the interface. When you identify an interface and direction as a blocking interface, the IPS takes control of the ACL on that interface/direction.
	If you are managing the blocking device in Security Manager, you can identify the ACL name by selecting the blocking device, then selecting **Tools > Preview Config**. Look for the **ip access-group** command in the interface configuration, and check the direction. For example, the following lines show that there is an ACL named CSM_FW_ACL_GigabitEthernet0/1 in the In direction attached to the GigabitEthernet0/1 interface.
	`interface GigabitEthernet0/1`
	`  ip access-group CSM_FW_ACL_GigabitEthernet0/1 in`
	In this example, if you configure GigabitEthernet0/1 in the In direction as a blocking interface, ensure that you specify CSM_FW_ACL_GigabitEthernet0/1 as a pre- or post-ACL. In most cases, you should specify the ACL as the post-ACL, so that the relatively short IPS blocking ACL first filters out undesirable traffic before the blocking device implements your other access rules.

Cat6k Block VLAN Dialog Box

Use the Add or Modify Cat6k Block VLAN dialog box to configure a blocking VLAN on a Catalyst 6500/7600 device that runs the Catalyst operating system and that is configured as an IPS blocking device. The IPS sensor uses the VLAN for blocking actions.

Tip If the Catalyst 6500/7600 runs Cisco IOS Software, add the device as a router, not a Cat6K.

Navigation Path

From the Add or Modify Cat6K Device dialog box, click the **Add Row** button beneath the VLANs table, or select a row in the table and click the **Edit Row** button. For information on opening the Cat6K Device dialog box, see Router, Firewall, Cat6K Device Dialog Box, page 43-14.

Field Reference

Table 43-7 Cat6k Block VLAN Dialog Box

Element	Description
VLAN	The number of the VLAN on the Catalyst 6500/7600 device that the IPS should use for blocking. The number can be 1 to 4094 and must be defined on the device.
Pre VACL Name Post VACL Name	The VLAN ACLs to combine with the blocking entries that the IPS creates to implement blocking actions. The Pre VACL is added before the blocking VACL, and the Post VACL is added after the blocking VACL. For more information, see Understanding Router and Switch Blocking Devices, page 43-4. **Tip** If you have configured a VACL on the VLAN, you must specify the name of the VACL in the Pre or Post VACL Name field or the VACL will be removed from the VLAN. When you identify a VLAN as a blocking interface, the IPS takes control of the VACL on that VLAN. Typically, you would specify the VACL name as the post-VACL.

Never Block Host or Network Dialog Boxes

Use the Add or Modify Never Block Host or Network dialog boxes to specify a host or network that should never be subject to blocking. The name of the dialog box indicates whether you are adding a host or network address.

Enter the IP address or the name of a network/host policy object that specifies the address. You can also click **Select** to select an object from a list or to create a new object. When selecting objects, the object can contain a single entry of the appropriate type. Host addresses do not have subnet masks (for example, 10.100.10.1), whereas network addresses have masks (for example, 10.100.10.0/24).

Navigation Path

From the IPS Blocking policy, select the Never Block Hosts or Networks tab and click the **Add Row** button or select an existing row and click the **Edit Row** button. Hosts and networks are listed in separate tables, so ensure that you click the buttons associated with the desired table. For information on opening the Blocking policy, see Blocking Page, page 43-8.

Managing IPS Sensors

Note From 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements as IPS is now End of Life. For more information, see EOL notice.

To perform day-to-day sensor management, you typically need to use a device manager such as the IPS Device Manager (IDM). Security Manager is focused on policy and event management.

However, the following topics describe some management activities that you can perform using Security Manager:

- Managing IPS Licenses, page 44-1
- Managing IPS Updates, page 44-4
- Managing IPS Certificates, page 44-10
- Rebooting IPS Sensors, page 44-12

Managing IPS Licenses

The following topics explain how to manage licenses for IPS devices:

- Updating IPS License Files, page 44-1
- Redeploying IPS License Files, page 44-2
- Automating IPS License File Updates, page 44-3

Updating IPS License Files

You can use Security Manager to update the licenses for IPS devices. This procedure explains how to update the licenses manually by retrieving them from Cisco.com or from a license file on the Security Manager server. For information about setting up automatic license updates, see Automating IPS License File Updates, page 44-3.

Before You Begin

If you use Cisco.com, you must first configure the IPS Update server to be Cisco.com, so that you can specify the username and password. You must use Cisco.com for licensing if you are using a device that requires it; for example, an IPS 4270 or an AIP SSM-40 in an ASA device requires a Cisco.com account. For information on configuring Cisco.com as the IPS Update server, see Configuring the IPS Update Server, page 44-4.

If you use local licenses, you must download them directly to the Security Manager server file system. You cannot do this through Security Manager; you must log into Windows on the server to download the licenses.

Related Topics

- Redeploying IPS License Files, page 44-2

Step 1 Select **Tools > Security Manager Administration** and select **Licensing** from the table of contents.

Step 2 Click the **IPS** tab (see IPS Tab, Licensing Page, page 11-58).

The table lists all IPS devices in the device inventory and displays the status of their licenses. The status can be valid, invalid, expired, no license, or trial license. The expiration date for the license is also shown. Click **Refresh License** to update the table with the latest license information from the devices (you can select one or more devices to limit the scope of the refresh).

To update licenses, do one of the following:

- To update devices with licenses obtained directly from Cisco.com—Select the devices you want to update and click **Update Selected via CCO**. A dialog box opens that lists the devices that can be updated from Cisco.com, which might not be all of the devices you selected. Review the list and click **OK**. The status of the update task is shown in the License Update Status Details dialog box (see License Update Status Details Dialog Box, page 11-62).

  To successfully update the license using this method, you must have a Cisco.com support contract that includes the serial numbers of the selected devices.

Tip The Cisco software license server (SWIFT) that contains the licenses might block requests from the same server for more than 9 licenses within a three minute period. Thus, you should select fewer than 9 devices at a time when performing manual license updates.

- To update devices with licenses that you have copied to the Security Manager server—Click **Update from License File**. A dialog box opens where you can select the license files. Click **Browse** to select them from the Security Manager local file system. You can select more than one license file. When you have selected the desired files, click **OK** to have them applied to the devices.

Redeploying IPS License Files

If an attempt to apply an IPS license update to a device fails, you can redeploy the update. Redeployment works only if you have already attempted to apply an update and a license file is associated with the IPS device.

Related Topics

- Updating IPS License Files, page 44-1

- Automating IPS License File Updates, page 44-3

Step 1 Select **Tools > Security Manager Administration** and select **Licensing** from the table of contents.

Step 2 Click the **IPS** tab (see IPS Tab, Licensing Page, page 11-58).

Step 3 Select the devices to which you want to redeploy licenses and click **Redeploy Selected Licenses**. A dialog box opens listing devices whose licenses you are redeploying. Click **OK** to perform the update.

The status of the update task is shown in the License Update Status Details dialog box (see License Update Status Details Dialog Box, page 11-62).

Automating IPS License File Updates

Security Manager can automatically apply IPS license updates to your IPS devices on a regular schedule. To successfully configure automatic updates, you must have a Cisco.com support contract that includes the serial numbers of your IPS devices.

Tip Security Manager applies new licenses only if the downloaded license has an expiration date further into the future than the one replaced or if the license information is different.

Before You Begin

You must first configure the IPS Update server to be Cisco.com, so that you can specify the Cisco.com username and password. For information on configuring Cisco.com as the IPS Update server, see Configuring the IPS Update Server, page 44-4.

Related Topics

- Updating IPS License Files, page 44-1
- Redeploying IPS License Files, page 44-2

Step 1 Select **Tools > Security Manager Administration** and select **Licensing** from the table of contents.

Step 2 Click the **IPS** tab (see IPS Tab, Licensing Page, page 11-58).

Step 3 Select **Download and apply licenses** and configure the following settings:

- **Days before the expiration date**—Select the number of days before a license expires that Security Manager should download an updated license. The default is 1 day.
- **Discover devices daily at**—Select the time of day when Security Manager should download licenses. At this time, Security Manager will check the license status on the devices, and contact Cisco.com for new licenses for devices that have no license, have expired licenses, or that have licenses that will expire within the number of days you selected.
- **Email License Update Results**—Select whether Security Manager should send e-mail notification of license update results. E-mails are sent with license expiration status and for license update job results. If you select this option, enter one or more e-mail addresses in the **Email Notification** field. Separate multiple addresses with commas.

  For the e-mails to be sent, you must configure an SMTP server as described in Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.

Step 4 Click **Save** to save your changes.

Managing IPS Updates

You can use Security Manager to apply sensor and signature updates to your IPS devices and shared policies. Through Security Manager, you can download updates and either set up automatic updates or apply them manually.

Signature updates are available only for IPS 5.1(4) and higher.

Tip If you have problems applying patches, service packs, or signature updates, check the time on your IPS sensor. If the time on the sensor is ahead of the time on the associated certificate, the certificate is rejected and the update may fail. Use the Network Time Protocol (NTP) to maintain accurate time on an IPS sensor. For information on configuring NTP on the sensor, see Identifying an NTP Server, page 36-23.

The IPS packages included with Security Manager do not include the package files that are required for updating IPS devices. You must download IPS packages from Cisco.com or your local update server before you can apply any updates. The downloaded versions include all required package files and replace the partial files that are included in the Security Manager initial installation.

The following topics describe how to use Security Manager to manage IPS updates:

- Configuring the IPS Update Server, page 44-4
- Checking for IPS Updates and Downloading Them, page 44-5
- Automating IPS Updates, page 44-6
- Manually Applying IPS Updates, page 44-7

Configuring the IPS Update Server

To apply IPS sensor and signature updates, Security Manager must download the updates to the Security Manager server from an identified IPS Update server.

You can use Cisco.com as the IPS Update server. Using Cisco.com ensures that the latest updates are available to you at their earliest availability. However, if you cannot use Cisco.com for some reason, you can set up your own local IPS Update web server, manually download updates to it, and configure Security Manager to obtain the updates from your local server.

Tip If you are using a device that requires a Cisco.com login for updating licenses, such as an IPS 4270 or an AIP SSM-40 in an ASA device, you must configure the IPS Update server as Cisco.com. You cannot use a local server.

Related Topics

- Automating IPS Updates, page 44-6
- Manually Applying IPS Updates, page 44-7

Step 1 Select **Tools > Security Manager Administration** and select **IPS Updates** from the table of contents to open the IPS Updates page (see IPS Updates Page, page 11-47).

Step 2 In the Update Server area, click **Edit Settings** to open the Edit Update Server Settings dialog box (see Edit Update Server Settings Dialog Box, page 11-52).

Step 3 Enter the identifying information for your server. Based on the server type selected in the Update From field:

- Cisco.com—Enter a Cisco.com username and password. The user account you specify must have applied for eligibility to download strong encryption software. To verify the account has the appropriate permissions, go to Cisco.com and try to download an IPS update package. You will be prompted to accept the appropriate agreements if the account is not already qualified.
- Local server—Enter the IP address or DNS host name of your server, a username and password if you require a log in before allowing access, and the path to the folder that contains the files. For the path, do not enter the entire URL; enter only the path portion of the URL (for example, the path in http://servername/IPSpath is IPSpath). Also add IIS configuration settings:
  - Home Directory should have listing enabled.
  - Documents should have Default Content Page disabled.

Enter certificate information. Before you can download an IPS package, you must accept the Cisco.com certificate. You must accept the certificate from both the "Image Meta-data locator" site and the download site of the IPS packages to start downloading images successfully (see Edit Update Server Settings Dialog Box, page 11-52).

If your network requires a proxy server to get from the Security Manager server to the IPS Update server, select **Enable Proxy Server** and enter the information for the proxy server.

Click **OK** to save your changes.

Step 4 Click **Save** on the IPS Updates page. Your changes are not completely saved unless you click **Save**.

Step 5 Test the connectivity to the IPS Update server by clicking **Download Latest Updates**. A dialog box opens. Click **Start** to have Security Manager log into the update server, check for new updates, and download them. The dialog box displays the results of the operation.

If you are using Cisco.com and experience a download failure, double-check the user account to ensure it has the required permissions for downloading strong encryption software.

Checking for IPS Updates and Downloading Them

You can use Security Manager to check for IPS sensor and signature updates and download them to the Security Manager server, where you can apply them to your IPS devices and policies.

You can manually download IPS updates, automate IPS update downloads, or download them when you try to manually apply them to a device. The following procedure explains how to manually check for updates and download them. For information on configuring automatic downloads, see Automating IPS Updates, page 44-6. For information on downloading updates while manually applying them to devices or policies, see Manually Applying IPS Updates, page 44-7.

Before You Begin

You must configure the IPS Update server as described in Configuring the IPS Update Server, page 44-4.

Related Topics

- Automating IPS Updates, page 44-6
- Manually Applying IPS Updates, page 44-7

Step 1 Select **Tools > Security Manager Administration** and select **IPS Updates** from the table of contents to open the IPS Updates page (see IPS Updates Page, page 11-47).

Step 2 Review the status information in the Update Status group, and do any of the following:

- Click **Check for Updates**. A dialog box opens to display the results of the operation. Click **Start** to have Security Manager log into the IPS Update server and check for updates.
- Click **Download Latest Updates**. A dialog box opens to display the results of the operation. Click **Start** to have Security Manager log into the IPS Update server, check for updates, and download them to the Security Manager server.

Tip If a Cisco.com download fails, ensure that the account you are using has applied for eligibility to download strong encryption software. For details, see the description of User Name in Edit Update Server Settings Dialog Box, page 11-52.

Automating IPS Updates

You can automatically apply sensor image and signature updates to compatible IPS devices to ensure that they are up to date. If desired, you can partially automate the updates to maintain the desired level of control over the process.

Tip If you later decide that you did not want to apply a signature update, you can revert to the previous update level by selecting the Signatures policy on the device, clicking the **View Update Level** button, and clicking **Revert**.

Tip If you do not manage IPS devices, consider taking the following performance tuning step. In *$NMSROOT*\MDC\ips\etc\sensorupdate.properties, change the value of packageMonitorInterval from its initial default value of 30,000 milliseconds to a less-frequent value of 600,000 milliseconds. Taking this step will improve performance somewhat. [*$NMSROOT* is the full pathname of the Common Services installation directory (the default is C:\Program Files\CSCOpx).]

Before You Begin

You must configure the IPS Update server as described in Configuring the IPS Update Server, page 44-4.

Related Topics

- Checking for IPS Updates and Downloading Them, page 44-5
- Manually Applying IPS Updates, page 44-7
- Understanding IPS Network Sensing, page 36-2
- Deploying Configurations in Non-Workflow Mode, page 8-28

- Deploying Configurations in Workflow Mode, page 8-34

Step 1 Select **Tools > Security Manager Administration** and select **IPS Updates** from the table of contents to open the IPS Updates page (see IPS Updates Page, page 11-47).

Step 2 In the Auto Update Settings group in the lower portion of the page, select an auto update mode to establish the extent of automation. Choices include:

- Download, Apply, and Deploy Updates—Security Manager checks for updates according to your schedule, downloads them to the Security Manager server, applies them to the selected devices and policies, and starts a deployment job to update the affected devices. This choice ensures that your devices are running the latest updates with minimal effort for your operations staff.
- Disable Auto Update—Security Manager does not perform any automatic actions for IPS updates.
- Check for Updates—Security Manager checks for updates according to your schedule and updates the information in the Update Status group. No devices or policies are updated.
- Download Updates—Security Manager checks for updates according to your schedule and downloads any new updates to the Security Manager server.
- Download and Apply Updates—Security Manager checks for updates according to your schedule, downloads them, and applies them to the selected devices and policies. You must separately create a deployment job to deploy the changes to the affected devices.

Step 3 Click **Edit Update Schedule** to open a dialog box where you can specify the schedule for the operation. Select the starting date, enter the starting time in 24-hour format (hh:mm), and select whether the schedule should be by the hour, day, week, month, or a one-time event. Click **OK** to save the schedule.

Step 4 (Optional) Enter an e-mail address in the Notify Email field. Security Manager will notify this user when a package is available for download or has been downloaded, applied, or deployed. You can enter more than one address by separating the addresses with commas.

Step 5 Select the devices and shared policies you want to automatically update in the Apply Update To selector. Use the Type field to toggle between local policies (for devices) and shared policies.

To select a device or policy, click it in the selector and click the **Edit Row** button (the pencil icon below the selector). This action opens the Edit Auto Update Settings dialog box. Select the types of updates you want to apply: minor sensor updates and service packs or service packs only, and the signature update level. Click **OK** to save your changes. The devices to which the policy apply are added to the Devices to be Auto Updated list. A message will indicate if you need to submit your changes for the change to take effect.

Step 6 Click **Save**.

Manually Applying IPS Updates

You can manually apply image and signature updates to compatible IPS devices using the Apply IPS Update wizard. Use this procedure with policies and devices that you did not configure for automatic updates (as described in Automating IPS Updates, page 44-6).

When applying signature updates, the wizard displays those signatures in the update that are not configured on the target IPS devices. You can configure the new signatures before they are applied.

When applying image and signature updates, only those devices to which the updates can be applied are available for selection. Inapplicable devices are grayed out. If you hover the mouse pointer over a grayed out device, a tooltip displays the reason for graying out the device. A device can be grayed out even if a

signature update applies to it but if the required engine upgrade or generic packages are not available. Following are some of the instances when devices might be grayed out and the corresponding tooltip labels:

- If the version of the selected signature or sensor package is lower than the version of the target IPS device, Security Manager grays out the device and a mouse-over tooltip displays the message "Selected package is inapplicable".
- If you try to use Security Manager to upgrade an IPS device from version 7.2.2, with SNMP policy configured, to version 7.3.1, a mouse-over tooltip displays the message "Selected upgrade is not recommended. Unassign the SNMP policy on the device and deploy it to continue with the upgrade to 7.3.1". This is because SNMPv3 is not supported in IPS version 7.3.1.
- If you try to use Security Manager to perform a signature update that does not contain one or more threat profiles applied to the device, Security Manager grays out the device and a mouse-over tooltip displays the message "Currently applied threat profile is not applicable to this signature version". It does not allow the signature update to be successfully applied. You must remove the existing threat profile and then proceed with the signature update.

Tip If you later decide that you did not want to apply a signature update, you can revert to the previous update level by selecting the Signatures policy on the device, clicking the **View Update Level** button, and clicking **Revert**.

Before You Begin

Configure the IPS Update server as described in Configuring the IPS Update Server, page 44-4.

Related Topics

- Checking for IPS Updates and Downloading Them, page 44-5
- Selecting a Signature Category for Cisco IOS IPS, page 45-6

Note This note describes a difference between the update packages for IPS 7.1.3 and those used for earlier versions. When you open the Apply IPS Update wizard (Tools > Apply IPS Update), the first page of the wizard lists the sensor and signature update packages that are available. Beginning with IPS 7.1.3, a single update package is used for all supported platforms, such as IPS-4270 and ASA-SSE-AIP-85; example: IPS-CSM-K9-7.1.3.zip. Prior to IPS 7.1.3, a separate package was used for each supported platform; example: IPS-CS-MGR-SSC_5-K9-6.2-4-E4.zip.

Step 1 Select **Tools > Apply IPS Update** to open the Apply IPS Update wizard.

Step 2 On the first page of the wizard, select the update that you want to apply. This page lists the sensor and signature updates that are available. Do the following on this page:

- To update the list of packages, click **Download Latest Updates**. Security Manager logs into the IPS Update server and downloads the updates that have become available since the last download. This works only if you have configured an update server as described in Configuring the IPS Update Server, page 44-4. You can also update the list of packages by doing the following:
  - Configure automatic downloads on the IPS Updates page (select **Tools > Security Manager Administration > IPS Updates**). For more information, see IPS Updates Page, page 11-47.
  - Manually download the updates to the CSCOpx\MDC\ips\updates folder in the product installation folder (typically Program Files) on the Security Manager server.

You can also check for updates without downloading them by clicking **Check for Updates**. The Update Status information is the same as described in IPS Updates Page, page 11-47.

- Select the signature or sensor update you want to apply to your IPS devices in the Updates Downloaded table. Use the **Type** field to toggle between the types of updates (you can select only one update to apply):
  - **Sensor Updates**—Displays the filename, the major, minor, service pack, and patch versions, as well as the supported engine release. You must apply all major sensor updates, however, minor updates are cumulative.
  - **Signature Updates**—Displays the filename, the signature number, and the supported engine release. Signature updates are cumulative; however, applying them as separate packages allows you to separate your work into more manageable units if you intend to tune the updates to match the specific needs of your network.

Note The engine package is not listed on the update page, but Security Manager implicitly pushes the engine package automatically in the case of a signature update that requires a higher engine version. (This occurs only when updating a device with the particular version that the engine package requires.)

Click **Next** to continue.

Step 3 On the second page of the wizard, select the local signature policies (representing devices not assigned to any shared signature policy) and shared signature policies you want to update from the Apply Updates To list. Use the **Type** field to toggle between the types of policies. You can select any combination of local and shared policies. When you select a policy, the devices that use the policy are selected for update.

To select all applicable devices or shared policies, click **Select All**. To erase your selection and start over, click **Deselect All**. These buttons apply only to the displayed list.

IPS devices to which the update does not apply are grayed out in the Apply Updates To list, and you cannot select them. When you select a device that can be updated, it is listed in the Devices Assigned to Selected Policies list; these are the only devices that will be updated. If you select a shared policy, all devices that are using the policy appear in the selected policies list, but the devices to which the update does not apply are grayed out.

Tip The engine release controls which devices you can select for sensor updates; you can apply the update only to devices that use the same engine version, regardless of the release version. For example, if your device is running 6.0(5) E3, you can update to 6.1(1) E3 but not to 6.1(1) E2. You also cannot apply a 6.1(1) E3 update to a device running 6.1(1) E2. If you want to update the engine version, select a signature update with the higher engine version, and Security Manager will update the engine level automatically while updating the signatures. For example, if the device has the 6.1(1) E2 version and needs to have the E3 engine package applied, choose the signature package that requires the E3 engine and apply it to the device; doing so applies the engine package automatically to the device while updating the signatures. Thus, if the device you want to update is grayed out, click **Back** and change your update selection.

If you are applying a signature update, and you want to edit the signatures before applying them, click **Next** to continue. Otherwise, click **Finish** to apply your update to the policies.

Step 4 (Optional) On the third page of the wizard, modify the signatures as desired.

The signatures list displays the new and modified signatures between the signature level of the selected update and the lowest signature level among the selected devices. If the selected devices include both IPS sensors and Cisco IOS IPS devices, the signatures for these devices appear on separate tabs.

Click the link in the ID number to read the description for the signature on Cisco.com. The Status column indicates whether the signature is new or modified (see the visual description of the icons on the wizard page).

To edit a signature, select it in the table and click the Edit button below the table (the pencil icon). For help in understanding the signature, click **Help** in the dialog box that the Edit button opens.

For details on available signature information, see Signatures Page, page 39-4. In the Signature Summary Table, you can also add custom signatures and delete signatures, but you cannot do that on this page of the Apply IPS Update Wizard.

Click **Finish** to apply your update to the policies and to save your edits.

Step 5 Submit your changes and deploy them to the devices. For information on creating deployment jobs, see these topics:

- Deploying Configurations in Non-Workflow Mode, page 8-28
- Deploying Configurations in Workflow Mode, page 8-34

Managing IPS Certificates

When you configure Security Manager to use SSL (HTTPS) to communicate with your IPS devices, the certificate configured on the device must match the certificate stored in Security Manager's certificate store. Mismatched certificates will result in communication failures during policy discovery or deployment.

IPS devices use self-signed certificates that have a fixed validity period of about 2 years. When the certificate expires, you need to regenerate the certificate and update the certificate store with the new certificate.

Security Manager includes a utility that you can use to synchronize the certificate store with the certificate defined on the device, to regenerate expired certificates, and to view the status of certificates (including expiration dates) on the IPS devices that you manage.

Tip If you are using HTTP for communication with the IPS devices, certificates are not used and you cannot manage them. IPS device communication settings are configured in the Security Manager Administration Device Communication page (see Device Communication Page, page 11-21).

The following procedure explains how to manage your IPS certificates with Security Manager.

Related Topics

- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48
- Manually Adding SSL Certificates for Devices that Use HTTPS Communications, page 9-5
- Security Certificate Rejected When Discovering Device, page 9-6
- Invalid Certificate Error During Device Discovery, page 9-7

Step 1 Select **Manage > IPS > IPS Certificates** to open the IPS Certificates dialog box.

Tip The list shown in this dialog box is not automatically refreshed. Click **Refresh** whenever you open the dialog box to ensure that you are looking at the most current certificate expiration information.

The dialog box lists all IPS sensors that are in the inventory according to their Security Manager display name. Not all columns are displayed (right-click any cell heading to select additional columns). The main columns of interest are the following:

- **Certificate Mismatch?**—Whether the certificate defined on the device is the same as that in Security Manager. This field is blank if the certificate is unavailable or non-retrievable; otherwise, it can have these values:
  - **No**—The device and Security Manager have the same certificate. No action is required.
  - **Yes**—The device and Security Manager have different certificates. If the certificate has not expired, select the device and click **Sync Certificates** to replace the certificate in the Security Manager certificate store with the certificate from the device.
- **Valid Until on Device, Valid From on Device**—These two separate columns show the date range within which the certificate is valid. The certificate expires after the Valid Until date is reached. Consider regenerating the certificate as this date approaches.
- **Certificate Status on Device**—Shows the current status of the certificate as it exists on the device:
  - **Valid Certificate**—The certificate is good and within the validity date range.
  - **Expired Certificate**—The certificate has passed its Valid Until date and is now expired. Select the device and click **Regenerate Certificate** to create a new valid certificate on the device and to have the certificate loaded into the Security Manager certificate store.
  - **Certificate Not Yet Valid**—The certificate has not yet reached its Valid From date and cannot be used yet. This might indicate a mismatch between the time settings on the device and on the Security Manager server. Ensure that the time settings are the same (consider using an NTP server). Consider regenerating the certificate.
  - **Unavailable – Refresh to get Cert Info**—The certificate is not currently in the Security Manager certificate store. Click **Refresh** to have Security Manager retrieve the certificate from the device and load it into the certificate store.
  - **Nonretrievable – Cert Info not available**—Security Manager was not able to log into the device and retrieve the certificate, or you are using HTTP for communications. Select the device and click **Refresh**.

    If refresh does not resolve the problem, ensure that the device is operating normally (that it is not down), and then check the device properties to ensure that correct credentials are configured for access (see Viewing or Changing Device Properties, page 3-41). If credentials are not the problem, also check the Allowed Hosts policy configured on the device and ensure that the Security Manager server is included as an allowed host (see Identifying Allowed Hosts, page 36-7). You can also log into Windows on the Security Manager server and use ping to see if there is a route between the server and the IPS device.
- **Thumbprint on CSM, Thumbprint on Device**—These separate columns show the thumbprint for the certificate in the certificate store and on the device.

Step 2 Use any of the following buttons to perform the indicated actions. Except where indicated, if you do not select one or more devices before clicking the button, the action is performed on all listed devices, which can be time-consuming if there are a lot of IPS devices. You are warned before an operation is performed on all devices and given the option to stop it.

- **Sync Certificate**—Synchronize the certificate information in the Security Manager certificate store with the certificate on the device. The device certificate replaces the one in the certificate store.
- **Regenerate Certificate**—Generate a new certificate on the device and then load the new certificate into the certificate store.
- **Refresh**—Refresh the status information by having Security Manager contact the devices and retrieve certificate information, such as validity dates, and compare the certificate with the one in the certificate store. This action updates the Certificate Status on Device column and also determines whether there is a certificate mismatch.
- **Export**—Exports the entire certificates table to a comma-separated values (CSV) file. You cannot export less than the entire table. You are prompted for a file name and folder on the Security Manager server.

Rebooting IPS Sensors

You can reboot an IPS sensor from Security Manager.

To reboot the sensor, select it in Device view, right-click and select **Reboot Device**. You are asked to confirm that you want to reboot.

Security Manager does not provide status information on the reboot process.

Configuring IOS IPS Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS, it does not support any bug fixes or enhancements.

Some Cisco IOS routers, such as integrated services routers (ISRs), include native IPS capabilities based on IPS 5.1 software. You can configure some basic IPS inspection on these devices to supplement IPS sensor inspection or to support small networks.

This chapter contains the following topics:

- Understanding Cisco IOS IPS, page 45-1
- Overview of Cisco IOS IPS Configuration, page 45-4

Understanding Cisco IOS IPS

You can use Cisco Security Manager with the Cisco IOS Intrusion Prevention System (IOS IPS) to manage intrusion prevention on Cisco routers that use supported Cisco IOS Software releases 12.4(11)T2 and higher.

The Cisco IOS IPS acts as an in-line intrusion prevention sensor, watching packets and sessions as they flow through the router and scanning each packet to match any of the Cisco IOS IPS signatures. When it detects suspicious activity, it responds before network security can be compromised and logs the event through Cisco IOS syslog messages or Security Device Event Exchange (SDEE).

You can configure Cisco IOS IPS to choose the appropriate response to various threats. The Signature Event Action Processor (SEAP) can dynamically control actions that are to be taken by a signature event on the basis of parameters such as fidelity, severity, or target value rating. You can configure these actions in Security Manager through the Signatures and Event Actions policies.

When packets in a session match a signature, Cisco IOS IPS can take any of the following actions, as appropriate:

- Send an alarm to a syslog server or a centralized management interface.
- Drop the packet.
- Reset the connection.
- Deny traffic from the source IP address of the attacker for a specified amount of time.
- Deny traffic on the connection for which the signature was seen for a specified amount of time.

Cisco developed its Cisco IOS software-based intrusion-prevention capabilities and Cisco IOS Firewall with flexibility in mind, so that individual signatures could be disabled in case of false positives. Generally, it is preferable to enable both the firewall and Cisco IOS IPS to support network security policies. However, each of these features can be enabled independently and on different router interfaces.

For an overall understanding of the Cisco IOS IPS configuration process, see Overview of Cisco IOS IPS Configuration, page 45-4.

This section contains the following topics:

- Understanding IPS Subsystems and Support of IOS IPS Revisions, page 45-2
- Cisco IOS IPS Signature Scanning with Lightweight Signatures, page 45-2
- Router Configuration Files and Signature Event Action Processor (SEAP), page 45-3
- Cisco IOS IPS Limitations and Restrictions, page 45-3

Understanding IPS Subsystems and Support of IOS IPS Revisions

Cisco Security Manager automatically supports minor revisions of IOS IPS. To identify minor revisions that are supported, the IPS subsystem version is needed.

The IPS subsystem version is a version number used to keep track of Cisco IOS IPS feature changes. The subsystem number is show in the device properties (right-click the device and select **Device Properties**). You can also use the command **show subsys name ips** at a command line on the router that is running Cisco IOS IPS to show the detailed Cisco IOS IPS subsystem version. The 3.x subsystems are equivalent to IPS 5.x. For a list of the supported subsystems by Cisco IOS Software release, see the *Supported Devices and Software Versions for Cisco Security Manager* on Cisco.com for this release of Security Manager.

An IPS subsystem version is minor if the version difference is limited at postfix. For example, a revision from 3.0.1 to 3.0.2 is considered minor. For another example, 3.0.1 to 3.1.1 is also considered a minor version change. However, minor revisions that include new features are not automatically supported by Cisco Security Manager.

Cisco IOS IPS Signature Scanning with Lightweight Signatures

The addition of Cisco IOS IPS signature scanning with lightweight signatures in Cisco IOS Release 15.0(1)M is an enhancement to Cisco IOS IPS that allows loading of larger signatures sets, without consuming significant additional memory or reducing the memory consumed by an existing signature set, by loading equivalent lighter-weight signatures. These signatures are referred to as lightweight signatures.

Security Manager can discover and tune custom signatures with LWEs on ISRs and modular access routers. Security Manager supports the following features for signatures with LWEs on ISRs and modular access routers:

- New signature types
- Signature categories
- New default signature category recognition
- New engine update levels
- Licensing status—bypassed, expired, or not installed

Router Configuration Files and Signature Event Action Processor (SEAP)

As of Cisco IOS Release 12.4(11)T, signature definition files (SDFs) are no longer used by Cisco IOS IPS. Thus, you cannot not use the deprecated built-in signature sets, 128.sdf, 256.sdf, and attack-drop.sdf, with Security Manager.

Instead, routers access signature definition information through a directory that contains three configuration files—the default configuration, the delta configuration, and the SEAP configuration. You configure the location using the **IPS > General Settings** policy.

SEAP is the control unit responsible for coordinating the data flow of a signature event. It allows for advanced filtering and signature overrides on the basis of the Event Risk Rating (ERR) feedback. ERR is used to control the level in which a user chooses to take actions in an effort to minimize false positives.

Signatures once stored in NVRAM are now stored in the delta configuration file.

Cisco IOS IPS Limitations and Restrictions

Cisco IOS IPS routers do not support all the features that are supported by dedicated IPS sensor appliances and service modules. In addition, routers that support IOS IPS might not allocate as much memory to IPS functionality as an IPS sensor does. The following limitations and restrictions are important:

- When configuring an IOS IPS device, select only the signatures that you need. If you select all signatures that are available in Security Manager, you might exceed the memory available on the IOS IPS router and deployment can fail, the device might fail to load all of the signatures, or performance might be significantly degraded. If you encounter deployment failures, select a reduced set of signatures and then redeploy the configuration to the device.
- Security Manager-managed routers being configured to use IOS-IPS for the first time cannot use the auto-update process for signature updates. You must first update the router before you use the auto-update process. Follow these steps:
  a. Push an E3 signature, for example, S317.
  b. Push an intermediate signature, for example, S470.
  c. Push the first E4 signature, for example, S485.
  d. Push subsequent E4 signatures until you reach the desired level. Note that each delta should be less than 10 MB in size.

  After you have updated the router, you can use the auto-update process to update the signatures. The auto-update process will be successful as each incremental change will not exceed the memory available on the router. For information on configuring automatic updates, see Automating IPS Updates, page 44-6.
- Virtual sensors are not supported by IOS IPS.
- When using event action filters with an IOS IPS router, only a subset of IPS actions are available for removal from an event that meets the criteria of the event action filter. For more information on available event actions, see Filter Item Dialog Box, page 40-9 and Understanding IPS Event Actions, page 40-2.
- IOS IPS is based on IPS Software 5.1. Therefore, features introduced in later versions of IPS Software are typically not available in IOS IPS. For example, you cannot configure the following features:
  – Global correlation.

- Anomaly detection.
- OS identification in the event action network identification policy.

Overview of Cisco IOS IPS Configuration

There are a wide variety of devices on which you can configure the Intrusion Prevention System. From a configuration point-of-view, you can separate the devices into two groups: dedicated appliances and service modules (for routers, switches, and ASA devices) that run the full IPS software; and IPS-enabled routers running Cisco IOS Software 12.4(11)T and later (Cisco IOS IPS).

The following procedure is an overview of IPS configuration on a Cisco IOS IPS router. For dedicated IPS devices, including IPS service modules installed in a router, see Overview of IPS Configuration, page 36-5.

Cisco IOS IPS is a more limited feature meant for branch offices and small to medium sized networks, or to distribute IPS throughout a network. You typically cannot employ as many signatures in a Cisco IOS IPS router compared to a dedicated appliance. You also cannot configure advanced features such as global correlation, because Cisco IOS IPS is based on IPS Software version 5.1. When configuring Cisco IOS IPS devices, you are mostly configuring standard router policies, because the device is a router that is running a few IPS features. In comparison, the platform policies for IPS appliances and service modules are specific to IPS software.

Tip Before configuring Cisco IOS IPS, read *Cisco IOS Intrusion Prevention System Deployment Guide* on Cisco.com.

Step 1 Install and connect the device to your network. Install the device software and perform basic device configuration. Install the licenses required for all of the services running on the device. The amount of initial configuration that you perform influences what you will need to configure in Security Manager. For information about required basic settings, see:

- Setting Up SSL on Cisco IOS Routers, page 2-4
- Setting Up SSH, page 2-5
- Configuring Licenses on Cisco IOS Devices, page 2-10
- Initial Preparation of a Cisco IOS IPS Router, page 45-5
- Selecting a Signature Category for Cisco IOS IPS, page 45-6

Step 2 Add the device to the Security Manager device inventory (see Adding Devices to the Device Inventory, page 3-7). When you add the device be sure to make the following selections:

- When adding from Network or Export File, ensure that you select **IPS Policies** for policy discovery.
- When adding from Configuration File or by Manual Definition, ensure that you select **IPS** from the **Options** list, or the device will not be IPS-capable from Security Manager's point of view.

Step 3 Configure the IPS general settings to specify the location of the IPS files on the router. For more information, see Configuring General Settings for Cisco IOS IPS, page 45-7.

Step 4 Configure the IPS interface rules to enable IPS and to identify the interfaces on which traffic will be subject to IPS inspection. For more information, see Configuring IOS IPS Interface Rules, page 45-9.

Step 5 Configure IPS signatures and event actions. Event action policies are easier to configure than creating custom signatures, so try to use event action filters and overrides to modify signature behavior before trying to edit specific signatures. For more information, see the following topics:

- Chapter 40, "Configuring Event Action Rules"
- Configuring Signatures, page 39-4

Step 6 Maintain the device:

- Update and redeploy configurations as necessary.
- Apply updated signature and engine packages. For information about checking for updates, applying them, and setting up regular automated updates, see Managing IPS Updates, page 44-4.

Initial Preparation of a Cisco IOS IPS Router

Before you add a Cisco IOS IPS router to the Security Manager inventory, you need to perform some preparatory steps. The white paper Getting Started with Cisco IOS IPS with 5.x Format Signatures on Cisco.com provides a step-by-step explanation of a basic configuration. Although you could do some of the steps after adding the router to Security Manager, such as configuring interface rules, you should do at least the basic steps.

The following procedure explains the steps you are required to complete in the CLI. These steps are required because Security Manager either cannot complete them, or it is simply easier to do it in the CLI (as a one-time configuration). The white paper includes additional steps that you can complete in the CLI, and Security Manager can discover your configuration when you add the device to the inventory. The more you do in CLI, the less you will have to configure in Security Manager.

Tip You also must complete the basic router configuration steps as explained in Setting Up SSL on Cisco IOS Routers, page 2-4, Setting Up SSH, page 2-5, and Configuring Licenses on Cisco IOS Devices, page 2-10. The following steps apply to the IPS configuration only.

Step 1 Create a directory for IPS files on flash. For example, the following command creates a directory named ips:

```
router# mkdir ips
Create directory filename [ips]?
Created dir flash:ips
```

At this point, you can optionally configure the router to use this directory for IPS, or you can do it later in Security Manager (in the IPS > General Settings policy). Use the following commands to configure it in CLI:

```
router# configure terminal
router(config)# ip ips config location flash:ips
```

Step 2 Configure the Cisco IOS IPS crypto key. The crypto key is used to verify the digital signature for the main signature file (sigdef-default.xml) whose contents are signed by a Cisco private key to guarantee its authenticity and integrity at every release.

You can obtain the CLI required for the key from http://download-sj.cisco.com/cisco/ciscosecure/ids/sigup/5.0/ios/realm-cisco.pub.key.txt (login to Cisco.com is required).

Tip Configuring the key through the CLI is probably the easiest way to do it. Alternatively, you can configure it in Security Manager by assigning the IOS_IPS_PUBLIC_KEY pre-defined FlexConfig object to the router's FlexConfig policy. For more information about FlexConfigs, see Chapter 7, "Managing FlexConfigs".

a. Open the text file and copy its contents to the clipboard (select all text then press Ctrl+C).
b. If necessary, enter **configure terminal** at the router CLI prompt.
c. Paste the copied text file at the router prompt.
d. Exit configuration mode.
e. Enter the **show run** command to confirm that the key was correctly configured.

Step 3 Syslog is configured for IPS notifications by default. If you want to use SDEE for notifications, enable SDEE:

```
router# configure terminal
router(config)# ip ips notify sdee
```

Step 4 Select a signature category to compile. For detailed information, see Selecting a Signature Category for Cisco IOS IPS, page 45-6.

Selecting a Signature Category for Cisco IOS IPS

Cisco IPS appliances and Cisco IOS IPS with IPS 5.x format signatures operate with signature categories. All signatures are grouped into categories; the categories are hierarchical. An individual signature can belong to more than one category. Top-level categories help to define general types of signatures. Subcategories exist beneath each top-level signature category. (For a list of supported top-level categories, use your router CLI help (?) with the **category** command.)

Router memory and resource constraints prevent a router from loading all Cisco IOS IPS signatures. Thus, it is recommended that you load only a selected set of signatures that are defined by the categories. Because the categories are applied in a "top-down" order, you should first retire all signatures, followed by "unretiring" specific categories. Retiring signatures enables the router to load information for all signatures, but the router does not build the parallel scanning data structure.

Retired signatures are not scanned by Cisco IOS IPS, so they do not fire alarms. If a signature is irrelevant to your network or if you want to save router memory, you should retire signatures, as appropriate.

Security Manager does not manage the signature category command. You cannot configure it directly with a policy. However, you can configure the FlexConfig policy to include a FlexConfig object that configures the command. There is a pre-defined object, IOS_IPS_SIGNATURE_CATEGORY, that you can use. If you want to configure a different category than basic, make a copy of the object and edit it. For information on how to use FlexConfigs, see Chapter 7, "Managing FlexConfigs".

Tip If you do not use the **category** command to select a subset of IPS signatures that the device will attempt to compile, Security Manager will configure the category command to enable the IOS IPS Basic category to prevent the device resources from being overloaded. You can change the category manually on the

device to select another set of signatures to compile. We recommend that you configure the category before adding the device to Security Manager; however, this is not possible if you add the device through manual definition.

The following example shows how to first retire all signatures, then to configure the basic category and unretire the basic signatures:

```
Router> enable
Router# configure terminal
Router(config)# ip ips signature-category
Router(config-ips-category)# category all
Router(config-ips-category-action)# retired true
Router(config-ips-category-action)# exit
Router(config-ips-category)# category ios_ips basic
Router(config-ips-category-action)# retired false
Router(config-ips-category-action)# exit
```

Configuring General Settings for Cisco IOS IPS

Use the General Settings page to specify the global settings used for Cisco IOS IPS properties defined for a particular router. The default settings are appropriate for most situations; however, you must specify an IPS configuration file location. If storing the configuration file on the router, you must first create the directory as described in Initial Preparation of a Cisco IOS IPS Router, page 45-5.

Navigation Path

- (Device view) Select **IPS > General Settings** from the Policy selector.
- (Policy view) Select **IPS (Router) > General Settings**, then select an existing policy or create a new one.

Related Topics

- Overview of Cisco IOS IPS Configuration, page 45-4
- Understanding Cisco IOS IPS, page 45-1

Field Reference

Table 45-1 General Settings Page

Element	Description
Block Traffic when IPS engine is unavailable	Whether to block all inspected traffic if the IPS engine is not available, for example, when the signature engine is being built or if it fails to build. If you select this option, any traffic specified for inspection is dropped if IPS cannot process it (also known as fail-closed mode). Otherwise, traffic is allowed to pass in accordance with the other rules in place on the router (the default).

Table 45-1 General Settings Page (continued)

Element	Description
Apply Deny Action On	Where to apply ACL entries to drop traffic for Deny Attacker Inline or Deny Flow Inline events. Select one of the following values: • **Ingress Interface** (the default)—Enforce the deny action on the interface attached to the network from which the traffic originated. • **IPS enabled interfaces**—Enforce the deny action on the interface on which the triggered IPS rule is applied. Enabling this option causes IOS IPS to apply the ACLs directly to the IPS interfaces, and not to the interfaces that originally received the attack traffic. If the router is not performing load balancing, do not enable this setting. If the router is performing load balancing, we recommend that you enable this setting.
SDEE Properties	
Maximum Subscriptions	The maximum number of concurrent SDEE subscriptions allowed, in the range of 1-3. An SDEE subscription is a live feed of SDEE events. The default is 1.
Maximum Alerts	The maximum number of SDEE alerts that you want the router to store, in the range of 10-2000. Storing more alerts uses more router memory. The default is 200.
Maximum Messages	The maximum number of SDEE messages that you want the router to store, in the range of 10-500. Storing more messages uses more router memory. The default is 200.
IPS Config Location Properties	
IPS Config Location	The location where the router will save IOS IPS specific configuration files. These configuration files are automatically updated every time the IOS IPS configuration is changed or updated from Security Manager. When the router reboots, the IOS IPS configuration is retrieved and restored from these configuration files. To specify a location on the router, enter the name of the directory. The directory must already exist; Security Manager does not create it. For example, flash:ips. **Note** If the router has a LEFS-based file system, you will be unable to create a directory in router memory. In this case, flash: is used as the config location. To specify a location on a remote system, specify the protocol and path of the URL needed to reach the location. For example, if you want to save the config files to an HTTP server, then enter http://172.27.108.5/ips-cfg. Supported servers for saving the IOS IPS configuration files are: http://, https://, ftp://, rcp://, scp://, and tftp://.

Table 45-1 General Settings Page (continued)

Element	Description
Max retries	When storing configuration files on a remote system, how many times the router is to attempt to contact the remote system. The default is 1.
Timeout seconds between retries	When storing configuration files on a remote system, how long the router is to wait before attempting to contact the configuration location again. The default is 1.

Configuring IOS IPS Interface Rules

Note From version 4.17, though Cisco Security Manager continues to support IOS and IPS features/functionality, it does not support any bug fixes or enhancements.

Use the IPS Interface Rules policy to enable IPS inspection on Cisco IOS IPS routers and to specify the interfaces that will be subject to IPS inspection. You can identify a subset of the traffic on the interface that is subject to inspection by configuring an ACL and by specifying the traffic direction relative to the interface.

Related Topics

- Overview of Cisco IOS IPS Configuration, page 45-4
- Understanding Cisco IOS IPS, page 45-1

Step 1 Do one of the following to open the Interface Rules policy you want to modify:

- (Device view) Select **IPS > Interface Rules** from the Policy selector.
- (Policy view) Select **IPS (Router) > Interface Rules** from the Policy selector. Select an existing policy or create a new one.

The policy shows any existing interface rules, including the rule name, the name of the ACL that defines which traffic is inspected (if any), and the interface and traffic direction that is inspected. If no ACL is specified, all traffic on the interface in the specified direction is inspected.

Although the rules are numbered, the sequence of rules has no effect on IPS processing.

Step 2 Select **Enable IPS** to enable the deployment of IOS IPS configuration to the device.

If Enable IPS is unchecked, IPS rules are removed from all the router interfaces, which disables IPS. Also, no signature or event action policy will be deployed.

Step 3 Configure the interface rules. The rules identify the interfaces, and traffic direction on the interface, that will be inspected by IPS. The rules can optionally include an ACL to identify a subset of traffic for inspection.

- To add a rule, click the **Add Row (+)** button and fill in the Add IPS Rule dialog box. For detailed information, see IPS Rule Dialog Box, page 45-10.
- To edit a rule, select it and click the **Edit Row (pencil)** button.

- To delete a rule, select it and click the **Delete Row (trash can)** button.

IPS Rule Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Use the Add or Edit IPS Rule dialog box to identify the traffic flows to be inspected using the active signature policy.

Navigation Path

From the Interface Rules policy, click the **Add Row** button to add a new rule, or select a rule and click the **Edit Row** button. For information on opening the Interface Rules policy, see Configuring IOS IPS Interface Rules, page 45-9.

Field Reference

Table 45-2 Add or Edit IPS Rule Dialog Box

Element	Description
Rule Name	The unique name for this IPS rule. IPS rule names are not case sensitive. You cannot use a rule name that contains the same characters as another one previously defined but using a different case. For example MYRULE and MyRule are the same.
ACL Name	The name of the ACL policy object that defines which traffic should be subject to IPS inspection. If you do not specify an ACL, all traffic on the interface/direction pairs listed in the Interface Pairs table is subject to inspection. **Tip** If you create an ACL, permit entries identify traffic that is subject to inspection, whereas deny entries identify traffic that is exempt from inspection. Remember that there is an implicit deny any any rule at the end of the ACL, so if your intention is simply to identify exempt traffic, be sure to add a permit any any rule at the end of the ACL. Enter the name of the ACL policy object, or click **Select** to select it from a list or to create a new object.
Interface Pairs table	The interfaces and traffic direction pairs that are subject to IPS inspection. • To add a pair, click the **Add Row (+)** button and fill in the Adding Pair dialog box. See Pair Dialog Box, page 45-11. • To edit a pair, select it and click the **Edit Row (pencil)** button. • To delete a pair, select it and click the **Delete Row (trash can)** button.

Pair Dialog Box

Use the Adding or Editing Pair dialog box to identify the interface and traffic direction pair to add to a Cisco IOS IPS interface rule. For information on configuring interface rules, see Configuring IOS IPS Interface Rules, page 45-9.

Navigation Path

From the Add or Edit IPS Rule dialog box, click the **Add Row** button to add a new pair, or select a pair and click the **Edit Row** button. For information on opening the Add or Edit IPS Rule dialog box, see IPS Rule Dialog Box, page 45-10.

Field Reference

Table 45-3 Adding or Editing Pair Dialog Box

Element	Description
Direction	The traffic direction, with respect to the interface, on which IPS inspection should be performed. Select one of the following: • **In** (default)—The IPS rule should be applied to inbound traffic. • **Out**—The IPS rule should be applied to outbound traffic. • **Both**—The IPS rule should be applied to both inbound and outbound traffic.
Interfaces	The interface on which to apply this IPS rule. Enter the name of an interface or interface role object, or click **Select** to select the interface or interface role from a list or to create a new interface role. If you use interface roles, the rule is applied to all interfaces on the device that are defined by the role. The interfaces that match the role cannot conflict with an existing rule. You cannot specify the same interface for more than one interface rule.

PART 5

PIX/ASA/FWSM Device Configuration

Managing Firewall Devices

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches, PIX, FWSM, and IPS, it does not support any bug fixes or enhancements.

The following topics describe configuration and management of security services and policies on Cisco security devices: Adaptive Security Appliances (ASAs), PIX Firewalls, and the Catalyst 6500 series switch Services Modules—that is, Firewall Services Modules (FWSMs) and ASA-SMs.

This chapter contains the following topics:

- Firewall Device Types, page 46-1
- Default Firewall Configurations, page 46-2
- Configuring Firewall Device Interfaces, page 46-3
- VXLAN, page 46-73

Firewall Device Types

Security Manager can discover and manage a variety of Cisco security appliances or firewall devices, most notably the following:

- PIX 500 Series firewall devices
- ASA 5500 Series security appliances including the Cisco Virtual Security Appliance (ASAv)
- Security-specific Catalyst Services Modules

PIX 500 Series

The Private Internet eXchange (PIX) 500 Series firewall appliances are no longer sold, however they are still supported and a great many are still in use world-wide.

ASA 5500 Series

The Adaptive Security Appliance (ASA) 5500 Series devices provide comprehensive security services, including context-aware firewall capabilities and real-time threat defense. The ASA 5500 has replaced the PIX 500 as Cisco's primary security appliance. Visit the Cisco ASA 5500 Series Adaptive Security Appliance page on cisco.com for more information.

The Cisco ASAv Virtual Appliance, introduced in ASA 9.2(1), brings full firewall functionality to virtualized environments to secure data center traffic and multi-tenant environments. The ASAv runs on VMware vSphere. Although the ASAv is a virtual device, it is managed like other ASA devices in

Security Manager. For more information about the ASAv, see http://www.cisco.com/c/en/us/support/security/virtual-adaptive-security-appliance-firewall/tsd-products-support-series-home.html.

Note The ASAv does not support the following ASA features: Clustering, Multiple context mode, Active/Active failover, Ether channels, and Shared AnyConnect Premium Licenses.

Catalyst Services Modules

A variety of Services Modules (SMs) are available for the Catalyst 6500 switch, including two that provide firewall and security services. These are blade-type modules that are installed directly into the switch chassis.

The Firewall Services module (FWSM) allows any port on the switch to operate as a firewall port, integrating firewall security inside the network structure. See Cisco Catalyst 6500 Series Firewall Services Module for more information.

The Adaptive Security Appliance service module (ASA-SM) provides high-speed security services across Layers 2 through 7, and you can install up to four ASA-SM blades in a single switch, providing scalability to 64 Gbps. See Cisco Catalyst 6500 Series ASA Services Module for more information.

Note While the ASA-SM is a blade installed in a Catalyst 6500 switch—much like the FWSM physically—it is an ASA device, and it is documented as such. That is, refer to ASA-related topics for information about the ASA-SM. Where necessary, caveats and differences between the Service Module and the ASA appliance are noted.

Default Firewall Configurations

Firewall devices are shipped with certain settings already configured. When you manually add a newly installed firewall device to Cisco Security Manager, you should discover (import) the pre-set or default policies for that device. Importing these policies into Security Manager prevents them being unintentionally removed the first time you deploy a configuration to that device. For more information about importing policies, see Discovering Policies, page 5-12.

Cisco Security Manager provides a set of configuration files that contain default policies for a number of device types and versions. These configuration files are located in the directory: `<install_dir>`\CSCOpx\MDC\fwtools\pixplatform\ (for example, `C:\Program Files\CSCOpx\MDC\fwtools\pixplatform\`).

The file name indicates device type, operating system version, context support, and operation type. For example, “FactoryDefault_FWSM2_2_MR.cfg” is the configuration file for an FWSM, version 2.2, with support for Multiple contexts, operating in Routed mode. Similarly, “FactoryDefault_ASA7_0_1_ST.cfg” is the configuration file for an ASA, version 7.0.1, in Single-context, Transparent mode.

Refer to Interfaces in Single and Multiple Contexts, page 46-5 for more about security contexts, and Interfaces in Routed and Transparent Modes, page 46-5 for more about routed and transparent operation.

See Adding Devices from Configuration Files, page 3-22 for information about adding new devices from the supplied configuration files.

Configuring Firewall Device Interfaces

The Interfaces page displays configured physical interfaces, logical interfaces, and redundant interfaces, as well as hardware ports and bridge groups, for the selected device. From this page, you can add, edit and delete interfaces; enable communication between interfaces on the same security level; and manage VPDN groups and PPPoE users.

Note The Interfaces page displayed for ASA 5505 devices presents two tabbed panels: Hardware Ports and Interfaces. Similarly, the Interfaces page displayed for the Catalyst 6500 services modules (ASA-SMs and FWSMs) operating in transparent mode also presents two tabbed panels: Interfaces and Bridge Groups.

Navigation Path

To access the Interfaces page, select a security device in Device View and then select **Interfaces** from the Device Policy selector.

This section contains the following topics:

- Understanding Device Interfaces, page 46-3
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26
- Advanced Interface Settings (PIX/ASA/FWSM), page 46-68

Understanding Device Interfaces

An interface is a point of connection between a security device and some other network device. Interfaces are initially disabled; thus, as an essential part of firewall configuration, interfaces must be enabled and configured to allow appropriate packet inspection and forwarding.

There are two types of interface: physical and logical, where a physical interface is the actual slot on the device into which a network cable is plugged, and a logical interface is a virtual port assigned to a specific physical port. Generally, physical ports are referred to as interfaces, while logical ports are referred to as subinterfaces, virtual interfaces, VLANs, or EtherChannels, depending on their function. The number and type of interfaces you can define varies with appliance model and type of license purchased.

Note On devices running version 6.3 of the PIX operating system, the labels “physical” and “logical” are used, rather than “interface” and “subinterface.” Also, transparent mode and multiple contexts are not supported on these devices.

Subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs. Because VLANs keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or security appliances. This feature is particularly useful in multiple-context mode, allowing you to assign unique interfaces to each context.

As a general rule, interfaces attach to router-based networks, and subinterfaces attach to switch-based networks. All subinterfaces must be associated with a physical interface that is responsible for routing allowed traffic correctly.

If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. The physical interface must be enabled for the subinterface to pass traffic, but do not name the physical interface to ensure it does not pass traffic. However, if you do want to let the physical interface pass untagged packets, you can name the interface as usual. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for information about naming an interface.

Note The ASA 5505, combining switch and security appliance features, is a special case in that you configure both physical switch ports and logical VLAN interfaces. See Understanding ASA 5505 Ports and Interfaces, page 46-6 for more information.

The Catalyst 6500 services modules (ASA-SMs and FWSMs) do not include any external physical interfaces—instead, they use internal VLAN interfaces. For example, assume you assign VLAN 201 to an FWSM inside interface, and VLAN 200 to the outside interface. You assign these VLANs to physical switch ports, and hosts connect to those ports. When communication occurs between VLANs 201 and 200, the FWSM is the only available path between the VLANs, forcing traffic to be statefully inspected.

See the following sections for additional information about device interfaces:

- Interfaces in Routed and Transparent Modes, page 46-5
- Interfaces in Single and Multiple Contexts, page 46-5
- Understanding ASA 5505 Ports and Interfaces, page 46-6
- Configuring Subinterfaces (PIX/ASA), page 46-7
- Configuring Redundant Interfaces, page 46-8
- Configuring EtherChannels, page 46-9
- Configuring VNI Interfaces, page 46-15
- Configuring Tunnel Interface, page 46-22

Security Appliance Configurations

Firewall devices allow a variety of configurations, and the configuration determines how to define the interfaces associated with a specific device. The following table outlines the various configurations.

Table 46-1 Security Appliance Configurations

Device Type	Operational Mode (Router or Transparent)	Context Support (Single or Multiple)
PIX 6.3.x	N/A	N/A
PIX 7.0+/ASA	Router or Transparent	Single
PIX 7.0+/ASA, or security context of unmanaged PIX 7.0+/ASA	Router or Transparent	Multiple (see Checklist for Configuring Multiple Security Contexts, page 59-3)
FWSM, or security context of unmanaged switch (multiple mode)	Router or Transparent	Single or Multiple

Interfaces in Routed and Transparent Modes

Beginning with ASA/PIX 7.0 and FWSM 2.2.1, you can configure a security device to operate in one of two modes: *routed* or *transparent*. (The PIX 6.3 operates only in routed mode.)

In routed mode, the security appliance acts as a gateway or router for connected networks: it maintains IP addresses for its interfaces, and inspects and filters traffic traversing these interfaces based on IP address (Layer 3) information. In this mode, each device interface is connected to a different IP subnet, and has its own IP address on that subnet. Routed mode supports up to 256 interfaces in single mode or per context, with a maximum of 1000 interfaces divided between all contexts.

In transparent mode, the security appliance operates as a Layer 2 (data link) device, or transparent bridge, and is often referred to as a "bump in the wire," or a "stealth firewall." In this mode, you can define only two interfaces: inside and outside. The interfaces do not require IP addresses; they use VLAN IDs to forward inspected traffic. However, if the device includes a dedicated management interface, you can use it—either the physical interface or a subinterface—as a third interface for device-management traffic.

Note Cisco Security Manager does not populate the interface information for FWSM 2.x devices during discovery.

Bridge Groups

Beginning with the ASA 8.4.1 and FWSM 3.1, in transparent mode, you can increase the number of interfaces available to a device or context through use of bridge groups. You can configure up to eight bridge groups; on an FWSM each group can contain two interfaces; on an ASA 9.7.1 (Cisco Security Manager 4.13) each group can contain upto 64 interfaces. See Add/Edit Bridge Group Dialog Box, page 46-62 for more information.

Related Topics

- Chapter 47, "Configuring Bridging Policies on Firewall Devices"

Interfaces in Single and Multiple Contexts

Security "contexts" allow a single physical device to operate as multiple, independent firewalls. In multiple-context mode, each context defines a single virtual firewall, complete with its own configuration. Each context acts as a unique virtual firewall that inspects and filters traffic traversing the interfaces allocated to that context. Each context is "unaware" of other contexts defined on the same security appliance.

As with a single-context, routed-mode device, interfaces on a multiple-context device connect to router-based networks, subinterfaces connect to switch-based networks, and each subinterface must be associated with an interface that routes allowed traffic correctly.

However, you cannot define IP addresses, the routed-mode portion of the configuration, or identify the management interface until you have defined and deployed the contexts. But you cannot define a security context until you have defined the necessary interfaces and subinterfaces.

In other words, you must enable and configure the interfaces and subinterfaces on a device that will provide multiple security contexts (in either routed or transparent mode) before you can define and configure the security contexts themselves.

Refer to Chapter 59, "Configuring Security Contexts on Firewall Devices" for more information.

About Asymmetric Routing Groups

In some situations, return traffic for a session may be routed through a different interface than the one from which it originated. Similarly, in failover configurations, return traffic for a connection that originated on one unit may return through the peer unit. This most commonly occurs when two interfaces on a single FWSM, or two FWSMs in a failover pair, are connected to different service providers and the outbound connection does not use a NAT address. By default, the FWSM drops the return traffic because there is no connection information for that traffic.

You can prevent return traffic being dropped by assigning the VLAN interfaces on which this is likely to occur to an asymmetric routing (ASR) group. When a member interface receives a packet for which it has no session information, it checks the session information for other interfaces that are members of the same group.

If a match is not found, the packet is dropped. If a match is found, one of the following actions occurs:

- If the incoming traffic originated on a different interface on the same FWSM, some or all of the Layer 2 header is rewritten and the packet is re-injected into the stream.
- If the incoming traffic originated on a peer unit in a failover configuration, some or all of the Layer 2 header is rewritten and the packet is redirected to the other unit. This redirection continues as long as the session is active.

Note In failover configurations, you must enable Stateful Failover for session information to be passed from the standby unit or failover group to the active unit or failover group.

To assign an FWSM virtual interface to an asymmetric routing group, simply specify an ASR Group ID in the Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41. If the group does not exist, it is created and the interface assigned to it.

You must repeat the assignment for each interface that will participate in this ASR group. You can create up to 32 ASR groups and assign a maximum of eight interfaces to each group.

Note The upstream and downstream routers must use one MAC address per VLAN, and have different MAC addresses for different VLANs, to allow the redirection of packets from a standby unit to an active unit in failover configurations.

Understanding ASA 5505 Ports and Interfaces

The ASA 5505 is unique in that it includes a built-in switch, and there are two kinds of ports and interfaces that you need to configure:

- Physical switch ports – The ASA 5505 has eight Fast Ethernet switch ports that forward traffic at Layer 2, using the switching function in hardware. Two of these ports are power-over-Ethernet (PoE) ports. You can connect these ports directly to user equipment such as PCs, IP phones, or DSL modems. Or you can connect to another switch.
- Logical VLAN interfaces – In routed mode, these interfaces forward traffic between VLAN networks at Layer 3, using the configured security policy to apply firewall and VPN services. In transparent mode, these interfaces forward traffic between the VLANs on the same network at Layer 2, using the configured security policy to apply firewall services.

To segregate the switch ports into separate VLANs, you assign each switch port to a VLAN interface. Switch ports on the same VLAN can communicate with each other using hardware switching. But when a switch port on one VLAN attempts to communicate with a switch port on another VLAN, the ASA 5505 applies the security policy to the traffic, and routes or bridges between the two VLANs.

Note Subinterfaces and redundant interfaces are not available on the ASA 5505.

Navigation Path

The Interfaces page displayed for ASA 5505 devices presents two tabbed panels: *Hardware Ports* and *Interfaces*. To access these panels, select an ASA 5505 in Device View and then select **Interfaces** from the Device Policy selector.

Configuring ASA 5505 Switch Ports and Interfaces

Refer to Configuring Hardware Ports on an ASA 5505, page 46-61 for information about configuring the switch ports.

Refer to Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31 for information about configuring the interfaces.

Related Topics

- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26

Configuring Subinterfaces (PIX/ASA)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs. Because VLANs keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or security appliances. This feature is particularly useful in multiple-context mode, letting you assign unique interfaces to each context.

Note If you use subinterfaces, you typically do not also want the physical interface to pass traffic, as the physical interface passes untagged packets. Because the physical interface must be enabled for the subinterface to pass traffic, do not name the physical interface to ensure it does not pass traffic. However, if you do want to let the physical interface pass untagged packets, you can name the interface as usual.

Note This option is available only on PIX 7.0+ and non-5505 ASA devices.

Defining Subinterfaces

Follow these steps to configure a subinterface in the Add/Edit Interface (ASA/PIX 7.0+) dialog box, which is accessed from the device Interfaces page (see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26).

1. Choose **Subinterface** as the interface **Type** in the Add/Edit Interface dialog box.

The VLAN ID and Subinterface ID fields appear below the Hardware Port, Name and Security Level fields.

2. Choose the desired **Hardware Port** from the list of previously defined interface ports. If you do not see a desired interface ID, be sure that Interface is defined and enabled.
3. **VLAN ID** – Provide a VLAN ID for this subinterface: enter a value between 1 and 4094. The specified VLAN ID must not be in use on any connected device.

   Some VLAN IDs might be reserved on connected switches; see the switch documentation for more information. In multiple-context mode, you can only set the VLAN ID in the system configuration.
4. **Secondary VLAN ID** – Provide a secondary VLAN ID value for this subinterface; this enables the ASA to map the packets that arrive on the ASA on the secondary VLAN to a primary VLAN. configure: Enter a value between 1 and 4090.
   The secondary VLAN ID must be unique and not be the same as a VLAN ID. A secondary VLAN is supported on devices running ASA 9.5.2 or later in single context, in routed or firewall mode or as an L2 cluster.

Note You can add multiple VLAN IDs, separated by a space or a comma. You can also specify a range of VLAN IDs for e.g. 56-78.

5. **Subinterface ID** – Provide an integer between 1 and 4294967293 as the Subinterface ID. The number of subinterfaces allowed depends on your platform.

   For subinterface port identification, this ID is appended to the chosen Hardware Port. For example, *GigabitEthernet0.4* represents the subinterface assigned an ID of 4, operating on the port GigabitEthernet0.

Note You cannot change the Subinterface ID after you set it.

6. Continue configuring this interface, as described in Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

Configuring Redundant Interfaces

Beginning with Security Manager 3.2.2, you can define logical “redundant” interfaces to increase security appliance reliability. A redundant interface is a specific pair of physical interfaces, with one designated as active (or primary) and the other as standby (or secondary). If the active interface fails, the standby interface becomes active and starts passing traffic. This feature is separate from device-level failover, but you can configure redundant interfaces as well as failover, if desired. You can configure up to eight redundant interface pairs.

A redundant interface functions as a single interface (inside, outside, etc.), with only one of the member pair active at any one time. This redundant interface is configured normally, with a unique interface name, security level and IP address. Note that each member interface must be of the same type (e.g., GigabitEthernet), and cannot have a name, security level, or IP address assigned. In fact, do not configure any options other than Duplex and Speed on the member interfaces.

The redundant interface uses the MAC address of the first physical interface that you specify. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. Alternatively, you can explicitly assign a MAC

address to the redundant interface; this address is then used regardless of the member interface MAC addresses. In either case, when the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

Note This option is available only on PIX 8.0+ and non-5505 ASA devices.

Defining Redundant Interfaces

Follow these steps to configure two physical interfaces as a single logical "redundant interface" in the Add/Edit Interface (ASA/PIX 7.0+) dialog box, which is accessed from the device Interfaces page (see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26).

1. Choose **Redundant** as the interface **Type** in the Add/Edit Interface dialog box.

   The Redundant ID, Primary Interface and Secondary Interface options appear.

2. Provide an identifier for this redundant interface in the **Redundant ID** field; valid IDs are the integers from 1 to 8.

3. **Primary Interface** – Choose the primary member of the redundant interface pair from this list of available interfaces. Available interfaces are presented by Hardware Port IDs, as named interfaces cannot be used for a redundant interface pair.

4. **Secondary Interface** – Choose the secondary member of the redundant interface pair from this list of available interfaces. Available interfaces are presented by Hardware Port IDs, as named interfaces cannot be used for a redundant interface pair.

   Note Member interfaces must be enabled and of the same type (e.g., GigabitEthernet), and cannot have a Name, IP Address, or Security Level assigned. In fact, do not configure any options other than Duplex and Speed on the member interfaces.

5. Continue configuring this interface, as described in Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

Configuring EtherChannels

Beginning with ASA 8.4.1, you can define logical EtherChannel interfaces. An EtherChannel, also called a port-channel interface, is a logical interface consisting of a bundle of individual Ethernet links (a channel group). This provides increased bandwidth and fault tolerance compared to the individual links.

An EtherChannel interface is configured and used in the same manner as a single physical interface. You can configure up to 48 EtherChannels, each of which consists of between one and eight active Fast Ethernet, Gigabit Ethernet, or Ten-Gigabit Ethernet ports. For ASA 9.2(1), the number of active interfaces increased to 16.

Note You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the ASA if they do not use the same physical interfaces.

EtherChannel MAC Addressing

All interfaces that are part of a channel group share the same MAC address. This makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. By default, the EtherChannel uses the MAC address of the lowest-numbered member interface as its MAC address.

Alternatively, you can manually configure a MAC address for the port-channel interface. We recommend doing so in case the channel interface membership changes. For example, if you remove the interface that provides the port-channel MAC address, the port-channel is assigned the MAC address of the next lowest numbered interface, causing traffic disruption. Manually assigning a unique MAC address to the EtherChannel interface prevents this disruption. (Note that in multiple-context mode, you can assign unique MAC addresses to interfaces assigned to an individual context, including EtherChannel interfaces.)

About Management Only EtherChannel Interfaces

You can specify an EtherChannel group as a management-only interface, but note the following caveats:

- Routed mode – You must explicitly configure the EtherChannel to be Management Only in the Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31. Any non-management interface added to the management-only port-channel is treated as a management port. If you add an interface already defined as management-only to the management-only group, that attribute is ignored on the physical interface. Similarly, you cannot designate an interface as management-only if it is already a member of a management-only port-channel.
- Transparent mode – In this mode, members of a management-only EtherChannel can themselves only be management-only ports. Thus, when a management-only member is added to a transparent-mode EtherChannel, the channel inherits the management-only designation, while the designation is removed from the member interface. Conversely, when such an interface is removed from the EtherChannel, the designation is restored on the individual interface.

Using an EtherChannel Interface as a Failover Link

If an EtherChannel interface is specified as a failover link, all state-sync traffic for that link will travel over a single physical interface. Should that physical interface fail, the state-sync traffic will then traverse another physical interface that is part of the EtherChannel aggregated link. If there are no remaining available physical interfaces in the EtherChannel link specified for failover, the ASA falls back to the redundant interface, if one is specified.

While an EtherChannel interface is being used as an active failover link, changes to that EtherChannel configuration are not allowed. You can change the EtherChannel configuration of that link only by disabling either the link or failover, as follows:

- Disable the EtherChannel link while the configuration changes are being made, and then reactivate it (failover will not occur while the link is disabled).
- Disable failover while the configuration changes are being made, and then re-enable it (failover will not occur in the interim).

Note As with any other type of interface assigned as a failover link, the EtherChannel interface cannot be named. Further, none of the EtherChannel's member interfaces can be named.

Defining EtherChannels on an ASA

Follow these steps to configure multiple physical interfaces as a single logical EtherChannel interface in the ASA Add Interface or Edit Interface dialog boxes, which are accessed from the device Interfaces page (see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26).

Step 1 Choose **EtherChannel** as the interface Type.

The EtherChannel ID and interface-selection options appear on the General panel of the dialog box; the Load Balancing, LACP Mode, and Active Physical Interfaces: Minimum and Maximum fields appear on the Advanced panel.

Step 2 Provide an identifier for this EtherChannel in the EtherChannel ID field; valid IDs are the integers from 1 to 48. This number is appended to "Port-channel" to identify the EtherChannel in the Interface column of the table on the device's Interfaces page.

Step 3 **Available Interfaces** – Specify the members of this port-channel group by select one or more interfaces in this list of available interfaces, and then click the **>>** button to add them to the member list on the right.

Note All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

You can assign up to 16 interfaces to a channel group. For ASA 9.2(1) and later, each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces and ASA versions earlier than 9.2(1), only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure. Alternatively, you can create a static EtherChannel by setting LACP Mode to On (on the Advanced panel, as described below), which means all interfaces in the group can pass traffic.

Note After assigning interfaces to this EtherChannel group, you can edit the LACP Port parameters for each member interface, as described in Editing LACP Parameters for an Interface Assigned to an EtherChannel, page 46-12.

Step 4 Click the **Advanced** tab to display that panel.

Step 5 Choose a **Load Balancing** option in the EtherChannel section. See About EtherChannel Load Balancing, page 46-13, for more information about this option.

Step 6 Select the desired **LACP Mode**; the default is Active, which means up to eight interfaces are active, while up to eight are in stand-by mode, as determined by the Minimum and Maximum values under Active Physical Interfaces.

If you select On, a static port-channel is created in which all member interfaces are all "on," meaning you can have up to 16 ports passing traffic, with no stand-by ports. When you select this option, the Mode for all interfaces assigned to this EtherChannel group is switched to On (if the Mode for each is not already On). See Editing LACP Parameters for an Interface Assigned to an EtherChannel, page 46-12, for more information about this mode.

Step 7 Specify the Minimum and Maximum number of Active Physical Interfaces for this EtherChannel.

As mentioned, an EtherChannel can consist of between 1 and 8 active links for ASA devices earlier than 9.2(1) or between 1 and 16 active links for ASA 9.2(1)+. Use these fields to indicate the minimum and maximum number of interfaces that can be active in this channel group at any given time. If your switch does not support 16 active interfaces, be sure to set the maximum to 8 or fewer.

Step 8 Continue configuring this interface, as described in Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

Note The EtherChannel **LACP System Priority** for this device is specified in the Advanced Interface Settings (PIX/ASA/FWSM), page 46-68 dialog box.

Editing LACP Parameters for an Interface Assigned to an EtherChannel

After assigning interfaces to an EtherChannel (port-channel) group, you can edit the LACP Port parameters for each member interface, as described here.

Note This feature is available only on ASA 8.4.1+ devices.

The Link Aggregation Control Protocol (LACP) directs aggregation of physical Fast Ethernet, Gigabit Ethernet, or Ten-Gigabit Ethernet interfaces into an EtherChannel group, and updating the remote partner device with current information after it finds a compatible set of ports and assigns a unique value called an "operational key" to the group. Note that operational key assignment is automatic; you cannot configure it.

Caution These LACP parameters are not available when the EtherChannel is assigned as a failover link.

LACP System Priority

Every LACP-enabled device has a unique system ID that is formed by combining a System Priority identifier and the system's MAC address. In certain situations, two EtherChannel-linked systems may need to change the operational key assigned to a set of ports to allow optimal aggregation. In such a situation, the system with higher priority is allowed to dynamically modify the operational key value assigned to the ports to achieve better aggregation. The system with the lower priority is not allowed to change the operational keys. The System Priority identifier is user-configurable, as described in Advanced Interface Settings (PIX/ASA/FWSM), page 46-68.

LACP Port Parameters

Port identification is provided by a unique number assigned to every group interface; this identifier is formed by combining a configurable Port Priority number and the port number assigned to the interface.

The port identifier provides port aggregation priority. Ports are considered for active use in an aggregation starting with the port that has highest aggregation priority in the system, and working down through an ordered list of port identifiers. The use of this port aggregation priority makes aggregation predictable and reproducible by selecting the links for aggregation in the same manner when all links are running LACP concurrently.

In addition, you can configure the priority of each port to administratively control the set of stand-by ports. For example, the port with the lowest priority will be considered last for group aggregation and will be become a stand-by port (assuming enough members are assigned to the group to allow stand-by ports).

Related Topics

- Configuring EtherChannels, page 46-9

Editing LACP Port Parameters for an Existing EtherChannel Interface

Follow these steps to edit an existing EtherChannel-assigned interface:

Step 1 In the table on the device's Interfaces page, select an interface that is a Member of a Port-channel group. (See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for information about accessing and using this table.)

Step 2 Click **Edit Row** to open the Edit Interface dialog box for that interface.

Only the Enable Interface check box, the LACP Port parameters, and the Description field can be altered.

Step 3 Edit the **LACP Port** parameters as necessary:

- **Priority** – This number is combined with the port number assigned to the interface to produce a unique port identification number. This value can be 1 to 65535, with higher numbers signifying lower priorities. The default is 32768. This parameter applies only when the port is in Active or Passive mode.
- **Mode** – Choose one of these LACP modes:
  - **Active** – In Active mode, a port initiates LACP exchanges with the partner device and periodically sends updates to the partner. Active LACP reflects the port's preference to participate in the protocol regardless of the partner's control mode.
  - **Passive** – A Passive-mode port does not initiate LACP exchanges, but upon receiving a request from the partner, the port will start exchanging LACP information with the partner. Passive mode is useful when it is not clear if the remote port supports LACP.

    Some devices may show undesired behavior when they do not have LACP enabled and they receive periodic LACP updates. However, for channeling to operate correctly, at least one port must be configured in Active mode.
  - **On** – Use this mode to configure a static port-channel in which all member interfaces are "on," with no stand-by ports. No negotiation takes place and most restrictions associated with the other two modes do not apply; for example, the speed and duplex settings do not have to be the same for all member ports, and all member ports remain Active. Note that the remote ports also must be On. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.
- **VSS or vPC Switch ID** – Identifies the Virtual Switching System (VSS) or Virtual Port Channel (vPC) switch ID to which the interface is connected.

Step 4 Continue editing this interface, as described in Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

About EtherChannel Load Balancing

Traffic in an EtherChannel is distributed across the individual bundled links in a deterministic fashion; however, the load is not necessarily balanced equally across all the links. Instead, frames are forwarded on a specific link as a result of a hashing algorithm. This algorithm uses a specific field or combination of fields in the packet header to produce a fixed Result Bundle Hash (RBH) value that indicates which link to use.

The algorithm can use one or a combination of the following packet-header fields to determine link assignment: source IP address, destination IP address, source MAC address, destination MAC address, TCP/UDP port numbers, or VLAN IDs. The field combination used by the algorithm is chosen from the **Load Balancing** list (on the Advanced tab of the ASA's Add Interface and Edit Interface dialog boxes); these options are described in the following section. For additional information, see Configuring EtherChannels, page 46-9.

For example, suppose source MAC address (*src-mac*) is the chosen field: when packets are forwarded to an EtherChannel, they are distributed across the ports in the channel based on the source MAC address of each incoming packet. Therefore, to provide load balancing, packets from different hosts use different ports in the channel, but packets from the same host use the same port in the channel (and the MAC addresses learned by the device do not change).

Similarly, with destination MAC address forwarding, when packets are forwarded to an EtherChannel, each packet is distributed across the ports in the channel based on the packet's destination host MAC address. Thus, packets to the same destination are forwarded over the same port, and packets to a different destination are sent on a different port in the channel.

Therefore, when choosing a load-balancing option, use the option that provides the greatest variety in your configuration. For example, if most of the traffic on a channel is going only to a single MAC address, choosing the destination MAC address results in most of the traffic always using the same link in the channel. Alternatively, using source addresses or IP addresses might result in better load balancing, while a method that uses the source and destination addresses along with UDP or TCP port numbers can distribute traffic much differently.

Note This option is available only on ASA 8.4.1+ devices.

Load Balancing Options

When defining a single logical EtherChannel interface in the ASA Add/Edit Interface dialog box, choose one of the following **Load Balancing** options (on the Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41) to specify the basis of load distribution:

- **dst-ip** – Load distribution is based on the destination-host IP address only; the source of the packets is not considered. Each packet with the same destination IP address is forwarded over the same link.
- **dst-ip-port** – Load distribution is based on the destination-host IP address and TCP/ UDP port. This option offers more granularity and a little more complexity than destination IP address alone.
- **dst-mac** – Load distribution is based on the destination host MAC address of incoming packets.
- **dst-port** – Distribution is based on the destination port; that is, a TCP or UDP port and not a physical interface.
- **src-dst-ip** – Distribution is based on source and destination IP addresses—source and destination IP addresses are paired for hash calculations. This method provides more granularity than destination IP address, for example: packets to the same destination can be forwarded over different links in a port-channel if they are coming from a different IP source.
- **src-dst-ip-port** – Distribution calculation considers source and destination IP addresses, and TCP/ UDP ports. Provides even greater granularity and distribution.
- **src-dst-mac** – Calculation is based on source and destination MAC address pairing.
- **src-dst-port** – Load distribution is based on source and destination TCP/UDP port.
- **src-ip** – Based on source host IP address only.
- **src-ip-port** – Source IP address and TCP/UDP port.
- **src-mac** – Source MAC address only.
- **src-port** – Source TCP/UDP port only.
- **vlan-dst-ip** – Destination IP address and VLAN ID pairing.
- **vlan-dst-ip-port** – Combination of destination IP address, TCP/UDP port, and VLAN ID.
- **vlan-only** – VLAN ID only.

- **vlan-src-dst-ip** – Source and destination IP address, and VLAN ID.
- **vlan-src-dst-ip-port** – Source and destination IP address, TCP/UDP port, and VLAN ID.
- **vlan-src-ip** – Source IP address and VLAN ID.
- **vlan-src-ip-port** – Source IP address, TCP/UDP port, and VLAN ID.

Configuring VNI Interfaces

VNI interfaces are similar to VLAN interfaces: they are virtual interfaces that keep network traffic separated on a given physical interface by using tagging. You apply your security policy directly to each VNI interface. All VNI interfaces are associated with the same VTEP interface.

To configure VXLAN, you must first Configuring VXLAN Policy, page 46-73 and then create a VNI interface and associate the configured VXLAN policy to the VNI interface.

When VNI Interface is the chosen Type in the Add Interface or Edit Interface dialog box, the dialog box presents three tabbed panels of options: General, Advanced and IPv6.The following sections describe how to configure VNI interfaces using the three tabbed panels:

- VXLAN, page 46-73
- VNI Interfaces—General Tab, page 46-15
- VNI Interfaces—Advanced Tab, page 46-16
- VNI Interfaces—IPv6 Tab, page 46-17

VNI Interfaces—General Tab

When VNI Interface is the chosen Type in the Add Interface or Edit Interface dialog box, the dialog box presents three tabbed panels of options: General, Advanced and IPv6. The options provided by the **General** panel are described in this section.

Navigation Path

You can access the General panel in the Add Interface and Edit Interface dialog boxes, which are accessed from the ASA Interfaces page, as described in Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.

Related Topics

- Configuring VNI Interfaces, page 46-15
- VNI Interfaces—Advanced Tab, page 46-16
- VNI Interfaces—IPv6 Tab, page 46-17

Field Reference

Table 46-2 General tab: Add/Edit Interface Dialog Box (ASA)

Element	Description
Enable Interface	Check this box to enable the VNI interface if not already enabled.
Name	Enter the Interface Name. The name is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.
Security Level	Enter the Security Level, between 0 (lowest) and 100 (highest).

Table 46-2 General tab: Add/Edit Interface Dialog Box (ASA) (continued)

Element	Description
VXLAN	
VNI ID	Enter the VNI ID, between 1 and 10000. This ID is just an internal interface identifier.
VNI Segment ID	Enter the VNI Segment ID, between 1 and 16777215. The segment ID is used for VXLAN tagging.
Multicast Group IP Address	(Single Mode) Enter the Multicast Group IP Address. If you do not set the multicast group for the VNI interface, the default group from the VTEP source interface configuration is used, if available. If you manually set a VTEP peer IP for the VTEP source interface, you cannot specify a multicast group for the VNI interface. Multicast is not supported in multiple context mode.
NVE Mapped to VTEP Interface	Check the NVE Mapped to VTEP Interface check box. This setting associates the VNI interface with the VTEP source interface.
IP Type	Select the IP Type from the available options.
Static IP	IP Address—(Routed Mode) In the IP Address area, configure an IPv4 address. To configure IPv6, click the IPv6 tab. Subnet Mask—Specify the Subnet Mask.
Use DHCP	DHCP Learned Route Metric—(Required)To assign an administrative distance to the learned route, enter a value between 1 and 255 in the DHCP Learned Route Metric field. If this field is left blank, the administrative distance for the learned routes is 1. Obtain Default Route using DHCP—(Optional) Select this option to obtain a default route from the DHCP server so that you do not need to configure a default static route. Enable Tracking for DHCP Learned Route—(Optional) If Obtain Default Route using DHCP is selected, you can select this option to enable route tracking via a specific Service Level Agreement (SLA) monitor. The following option becomes available: Tracked SLA Monitor—Required if Enable Tracking for DHCP Learned Route is selected. Enter or Select the name of the SLA monitor object that defines the route tracking (connectivity monitoring) to be applied to this interface.
Description	(Optional) Specify a description for the interface.

VNI Interfaces—Advanced Tab

When VNI Interface is the chosen Type in the Add Interface or Edit Interface dialog box, the dialog box presents three tabbed panels of options: General, Advanced and IPv6. The options provided by the **Advanced** panel are described in this section.

Navigation Path

You can access the Advanced tab in the Add Interface and Edit Interface dialog boxes, which are accessed from the ASA Interfaces page, as described in Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.

Related Topics

- Configuring VNI Interfaces, page 46-15
- VNI Interfaces—General Tab, page 46-15
- VNI Interfaces—IPv6 Tab, page 46-17

Field Reference

Table 46-3 Advanced tab: Add/Edit Interface Dialog Box (ASA)

Element	Description
Active MAC Address	Use the Active MAC Address field to manually assign a private MAC address to the interface
Standby MAC Address	The Standby MAC Address field can be used to set a standby MAC address for use with device-level failover.
Roles	All interface roles assigned to this interface are listed in this field. Role assignments are based on pattern matching between the Name given to this interface and all currently defined Interface Role objects in Cisco Security Manager. Interface role objects are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces.
DHCP Relay Servers	Enter the IP address or select a Networks/Hosts object representing the interface-specific DHCP server to which DHCP requests on this interface are relayed. Use a comma to separate multiple values. You can configure a maximum of 4 interface-specific DHCP relay servers and a maximum of 10 global and interface-specific DHCP relay servers combined.
DHCP Relay Trust Info (Option 82)	Specifies that you want to trust this DHCP client interface. You can configure interfaces as trusted interfaces to preserve DHCP Option 82.

VNI Interfaces—IPv6 Tab

When VNI Interface is the chosen Type in the Add Interface or Edit Interface dialog box, the dialog box presents three tabbed panels of options: General, Advanced and IPv6. The options provided by the **IPv6** panel are described in this section.

Navigation Path

You can access the IPv6 panel in the Add Interface and Edit Interface dialog boxes, which are accessed from the ASA Interfaces page, as described in Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.

Related Topics

- Configuring VNI Interfaces, page 46-15
- VNI Interfaces—General Tab, page 46-15
- VNI Interfaces—IPv6 Tab, page 46-17

Field Reference

Table 46-4 IPv6 tab: Add/Edit Interface Dialog Box (ASA)

Element	Description
Enable IPv6	Check this box to enable IPv6 and configure IPv6 addresses on this interface. You can deselect this option to disable IPv6 on the interface, but retain the configuration information.
Enforce EUI-64	When selected, use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link is enforced. When this option is enabled on an interface, the source addresses of IPv6 packets received on the interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the interface identifier in an IPv6 packet is not in the Modified EUI-64 format, the packet is dropped and the following system log message is generated: `%PIX|ASA-3-325003: EUI-64 source address check failed.` Address format verification is performed only when a flow is created. Packets from an existing flow are not checked. Additionally, address verification can be performed only for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address. The Modified EUI-64 format interface identifier is derived from the 48-bit link-layer (MAC) address by inserting the hex number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link-layer address. To ensure the chosen address is from a unique Ethernet MAC address, the next-to-lowest order bit in the high-order byte is inverted (universal/local bit) to indicate the uniqueness of the 48-bit address. For example, an interface with a MAC address of 00E0.B601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

Table 46-4 IPv6 tab: Add/Edit Interface Dialog Box (ASA) (continued)

Element	Description
DAD Attempts	To specify the number of consecutive neighbor solicitation messages that are sent on an interface during duplicate address detection (DAD), enter a number from 0 to 600 in this field. Entering 0 disables duplicate address detection on the interface. Entering 1 configures a single transmission without follow-up transmissions; this is the default. Duplicate address detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection uses neighbor solicitation messages to verify the uniqueness of unicast IPv6 addresses. When duplicate address detection identifies a duplicate address, the state of the address is set to DUPLICATE and the address is not used. If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface and an error message similar to the following is issued: `%PIX-4-DUPLICATE: Duplicate address FE80::1 on outside` If the duplicate address is a global address of the interface, the address is not used and an error message is issued, similar to that shown previously for a duplicate link-local address. All configuration commands associated with the duplicate address remain as-configured while the state of the address is set to DUPLICATE. If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address, and all other IPv6 address associated with the interface are regenerated (that is, duplicate address detection is performed only on the new link-local address).
NS Interval	The interval between IPv6 neighbor solicitation retransmissions, in milliseconds. Valid values range from 1000 to 3600000 milliseconds; the default value is 1000 milliseconds. **Note** This value is included in all IPv6 router advertisements sent out on this interface.
Reachable Time	The amount of time, in milliseconds, within which a remote IPv6 node is considered still reachable, after initial reachability was confirmed. Valid values range from 0 to 3600000 milliseconds, the default value is 0. When 0 is used for the value, the reachable time is set as undetermined—it is up to the receiving devices to set and track reachable time. A configured time enables detection of unavailable neighbors. A shorter time allows detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.
Managed Config Flag	Whether or not to set the flag "managed-config-flag" in the IPv6 router advertisement packet.

Table 46-4 IPv6 tab: Add/Edit Interface Dialog Box (ASA) (continued)

Element	Description
Other Config Flag	Whether or not to set the flag "other-config-flag" in the IPv6 router advertisement packet.
Enable RA	When checked, IPv6 router advertisement transmissions are enabled on the interface. The following options are enabled: • **RA Lifetime** – The "router lifetime" value specifies how long nodes on the local link should consider the security appliance as the default router on the link. Valid values range from 0 to 9000 seconds; the default is 1800 seconds. Entering 0 indicates that the security appliance should not be considered a default router on the selected interface. Any non-zero value should not be less than the following RA Interval value. **Note** This value is included in all IPv6 router advertisements sent out on this interface. • **RA Interval** – The interval between IPv6 router advertisement transmissions on this interface. Valid values range from 3 to 1800 seconds, (or from 500 to 1800000 milliseconds if the following RA Interval in Milliseconds option is checked); the default is 200 seconds. The interval between transmissions should be less than or equal to the RA Lifetime value if it is non-zero. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value. • **RA Interval in Milliseconds** – Checking this option indicates that the provided RA Interval value is in milliseconds, rather than seconds.

Table 46-4 IPv6 tab: Add/Edit Interface Dialog Box (ASA) (continued)

Element	Description
Interface IPv6 Addresses	The IPv6 addresses assigned to the interface are specified in this section of the dialog box. • **Link-Local Address** – To override the link-local address that is automatically generated for the interface, enter the desired IPv6 link-local address in this field. The link-local address is composed of the link-local prefix FE80::/64 and the interface ID in Modified EUI-64 format. For example, an interface with a MAC address of 00E0.B601.3B7A would have a link-local address of FE80::2E0:B6FF:FE01:3B7A. An error will occur if another host is using the specified address. • **Enable Address Auto-Configuration** – Select this option to enable automatic configuration of IPv6 addresses on the interface using stateless autoconfiguration. The addresses are configured based on the prefixes received in Router Advertisement (RA) messages. If a link-local address has not been configured, then one is automatically generated for this interface. An error occurs if another host is already using the generated link-local address. – **Trust the DHCP Servers for default gateway**– Select this radio button to install a default route from Router Advertisements that come from a trusted source - the directly-connected network. – **Ignore trust and accept router advertisements** – Select this radio button to install a default route from Router Advertisements that come from another network. • The table in this section displays the IPv6 addresses assigned to this interface. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Address for Interface Dialog Box, page 46-52.
Interface IPv6 Prefixes	Use the table in this section to configure which IPv6 prefixes (that is, the network portion of the IPv6 addresses) are included in IPv6 router advertisements. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Prefix Editor Dialog Box, page 46-54.
Interface IPv6 DHCP	

Configuring Tunnel Interface

Cisco Security Manager 4.13 supports route based VPN method for the Site-to-Site VPN. This support requires configuration of the static crypto map access list and mapping it to an interface. Due to this requirement, Large Enterprises and Virtual Private Clouds need to track all remote subnets and include them in the crypto map access list. To overcome this challenge, ASA 9.7.1 is enhanced to support the route based VPN method using the VTI (Virtual Tunnel Interface). Thus, beginning with Cisco Security Manager 4.13, you can define tunnel interface for the VPN and its associated IPSec policy.

VTI is supported only for Regular IPsec with Hub and Spoke, and Point to Point VPN topologies. VTI is not supported for other topologies like Full Mesh Topology, Extranet VPN Topology and RAVPN Policies.

In a multi- hub and multi- spoke scenario, for the tunnel interface to establish connectivity from one peer to another peer, ensure interface roles are applied to the hubs and spokes.

Note The Advanced tab and IPv6 tab options are not applicable for VTI.

The following sections describe how to configure tunnel interface:

- Tunnel—General Tab, page 46-22
- Configuring IPSec Policy for Tunnel Interface, page 46-24

Tunnel—General Tab

In the Add Interface or Edit Interface dialog box, when you select Tunnel from the Type drop-down, the dialog box displays three tabs: General, Advanced, and IPv6. This section describes the options provided by the **General** panel.

Navigation Path

You can access the General panel from the ASA Interfaces page, as described in Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.

Related Topics

- Configuring Tunnel Interface, page 46-22
- Configuring IPSec Policy for Tunnel Interface, page 46-24

Field Reference

Table 46-5 General tab: Add/Edit Interface Dialog Box (ASA)

Element	Description
Enable Interface	Check this box to enable the tunnel interface if not already enabled.
Name	Enter the Interface Name. The name is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value.

Table 46-5 General tab: Add/Edit Interface Dialog Box (ASA) (continued)

Element	Description
Tunnel Interface	
Tunnel ID	Enter the unique Tunnel ID, between 0 and 100. This ID is an internal interface identifier. The specified ID is mapped with the interface name. The Name and the ID pair must be unique. This field is mandatory for Regular IPSEC VTI VPN.
Source Interface	Enter the source interface to be used for creating the VTI, IP address shall be picked up from this interface. Click the **Select** button to choose the source interface from the available interfaces. For more information, see Selecting Objects for Policies, page 6-2. **Note** Tunnel source and destination pair must be unique.
Destination IP/Hostname	The tunnel destination IP address to be used for the VTI. Beginning with 4.14, Cisco Security Manager allows you to specify a Hostname as the destination IP. **Note** Tunnel source and destination pair has to be unique.
IPv4 Mode	Check the check box to pass IPv4 as the tunnel protection mode, currently, only IPSec is supported. IPv4 network would be encapsulated in the tunnel.
IPSec Profile	Enter the IPSec profile to be attached to the tunnel interface. A policy object must have been created in the Policy Object Manager. For information on creating Policy Object, refer Configuring IPSec Policy for Tunnel Interface, page 46-24. **Note** If you select IPSec profiles with different IKEV1 transform sets for the peers, Cisco Security Manager will create the tunnel interface, but the connectivity between the two peers will not be established. To choose the profile from the IPSec Object Selector dialog, click the **Select** button. For more information, see Selecting Objects for Policies, page 6-2. **Note** When you specify the policy, ensure tunnel name is entered. Cisco Security Manager displays error message when the **Name** field is blank.
IP Type	From the drop-down, select Static IP. • IP Address—(Routed Mode) In the IP Address area, configure an IPv4 address. To configure IPv6, click the IPv6 tab. • Subnet Mask—Specify the Subnet Mask.
Description	(Optional) Specify a description for the interface.

Establishing Regular IPSec VPN Tunnel

The following checkpoints (while configuring the Tunnel - Configuring Tunnel Interface, page 46-22) help you to successfully establish the Regular IPSec VPN Tunnel connectivity:

1. You must enter Tunnel ID value.
2. Source interface must be configured and must be reachable to its peer through ISP or routing.
3. You must enter the peer source interface IP address in the Destination IP field.
4. For the IPSec Profile field:
   a. Select the same IKEV1 transform set for both the peer devices.
   b. In Point-to-Point topology, either one of the peers must be the responder.
   c. In Hub and Spoke topology, select the Hub as the responder; select all spokes as initiators.
5. IPV4 mode must be configured for enabling the interesting traffic.
6. You must enter IP address to establish the VPN, dynamic IP address is not supported.
7. Select Static or BGP routing for enabling the interesting traffic. In case of firewall policy, VTI is supported only in static routing.

Note Cisco Security Manager displays error message if BGP/Static route is not configured properly for Point-to-Point topology, and Hub and Spoke topology with one hub and one spoke. For Multi hub/spoke scenario, the error message is not displayed.

Configuring IPSec Policy for Tunnel Interface

Use the IPsec Policy page to configure the IPsec policy used during IKE Phase 1 and IKE Phase 2 negotiations for Regular IPsec with Hub and Spoke and Point to Point VPN topologies.

Navigation Path

- Choose **Manage > Policy Objects** to open the Policy Object Manager. Under All Object Types, click **IPSec Profile**. To add a profile, click the **Add** button.

Field Reference

Table 46-6 IPsec Profile

Element	Description
Name	Name of the IPSec policy.
Description	Description for the policy.
IKE Version	Choose the relevant IKE version— IKEv1 or IKEv2. **Note** Beginning with 4.14, Cisco Security Manager supports IKEv2. However, at a time, you can choose only one version of IKE.

Table 46-6 IPsec Profile

Element	Description
IKEv1 Transform Sets	The IKEv1 transform sets to be used for your tunnel policy. Transform sets specify which authentication and encryption algorithms are used to secure the traffic in the tunnel. You can select up to 11 transform sets. For more information, see Understanding Transform Sets, page 26-20. Transform sets may use only tunnel mode IPsec operation. You can associate more than one IKEv1 Transform Sets. If more than one of your selected transform sets is supported by both peers, the transform set that provides the highest security is used. **Note** For the tunnel to be functional, IKEv1 transform set on both peers should be the same. Click **Select** to select the IPSec transform set policy objects to use in the topology. If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27. This field is not available for IKEv2.
IKEv2 IPSec Proposal (ASA 9.8(1) Onwards)	Click **Select** to select the IPSec proposals to be used for your tunnel policy. Cisco Security Manager allows you to select more than one proposals. <check on the VM> If the required object is not yet defined, you can click the **Create (+)** button beneath the available objects list in the selection dialog box to create a new one. For more information, see Configuring IPSec IKEv1 or IKEv2 Transform Set Policy Objects, page 26-27. This field is not available for IKEv1.
Trustpoint (ASA 9.8(1) Onwards)	Click **Select** to select the CA servers for issuing certificates to the participating IPSec network devices. The peers that are configured with this policy obtains digital certificates from the selected CA server. You can specify only one trustpoint. For IKEv1, when trustpoint is used for authentication, the initiator should have the trustpoint specified under IPSec profile's trustpoint configuration; and for responder, the trustpoint should be specified under tunnel-group CLI (similar to non-VTI configuration). For IKEv2, when trustpoint is used for authentication, the trustpoint CLI is specified under tunnel-group CLI for both initiator and responder.
Certificate Chain (ASA 9.8(1) Onwards)	Select the check box to enable sending of the certificate chain for authorization. A certificate chain includes the root CA certificate, identity certificate, and key pair.
Responder Only	Check this check box to set the peer that is associated with this policy acts as the responder. Ensure only one of the peer is configured with the responder only settings.

Table 46-6 IPsec Profile

Element	Description
Enable Perfect Forward Secrecy (PFS) Modulus Group	Whether to enable the use of Perfect Forward Secrecy (PFS) to generate and use a unique session key for each encrypted exchange. In IPsec negotiations, PFS ensures that each new cryptographic key is unrelated to any previous key. If you select this option, also select the Diffie-Hellman key derivation algorithm to use when generating the PFS session key in the **Modulus Group** list. For an explanation of the options, see Deciding Which Diffie-Hellman Modulus Group to Use, page 26-7. The following Modulus Group is not supported for IKEv1. Ensure you do not select them for IKEv1: • group14 • group19 • group20 • group21 • group24
Lifetime (Seconds) Lifetime (Kilobytes)	The global lifetime settings for the Crypto IPsec security association (SA). You can specify the IPSec lifetime in seconds, in kilobytes, or both. • Seconds—The number of seconds an SA will exist before expiring. Enter a value within the range 120-2147483647 seconds. • Kilobytes—The volume of traffic (in kilobytes) that can pass between IPsec peers using a given SA before it expires. Valid values depend on the device type. Enter a value within the range 10-2147483647. To allow unlimited You can select Enable Unlimited Lifetime (Kilobytes) check box.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device	Select to allow the properties of this object to be redefined on individual devices. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Managing Device Interfaces, Hardware Ports, and Bridge Groups

The Interfaces page displays the interfaces, subinterfaces, redundant interfaces, virtual interfaces (VLANs), and EtherChannel interfaces, as well as the hardware ports and bridge groups, configured on the selected device, and lets you add, edit and delete them.

The types of interface available depend on device type, operating system version, and mode (routed or transparent). For example, EtherChannel interfaces are available only on ASA 8.4.1 and later devices, in both routed and transparent mode. See Understanding Device Interfaces, page 46-3 for more information.

Note The Interfaces page displayed for ASA 5505 devices presents two tabbed panels: Interfaces and **Hardware Ports**. Similarly, the Interfaces page displayed for both Firewall Services Modules (FWSMs), version 3.1 and later, and ASAs version 8.4.1 and later, operating in transparent mode also present two tabbed panels: Interfaces and **Bridge Groups**. Links to configuration information for these features are included in the following procedure.

Each security device must be configured, and each active interface must be enabled. Inactive interfaces can be disabled. When disabled, the interface does not transmit or receive data, but its configuration information is retained.

If you bootstrapped a new security device, the set-up feature configures only the addresses and names associated with the inside interface. You must define the remaining interfaces on that device before you can specify access and translation rules for traffic traversing that security device.

Transparent firewall mode allows only two interfaces to pass traffic; however, if your platform includes a dedicated management interface, you can use it (either the physical interface or a subinterface) as a third interface for management traffic.

Follow these general steps to manage security-device interfaces and related options. You can add, edit and delete configured interfaces, subinterfaces, redundant interfaces, virtual interfaces (VLANs), EtherChannel interfaces, hardware ports, and bridge groups, according to the type of device selected.

Step 1 Ensure Device View is your present application view; if necessary, click the **Device View** button on the toolbar.

Note For more information on using the Device View to configure device policies, see Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30).

Step 2 Select the security device you want to configure.

Step 3 Select **Interfaces** in the Device Policy selector.

The Interfaces page is displayed. The information displayed, and the page itself, varies based on the selected device type and version, the operational mode (routed versus transparent), and whether the device hosts single or multiple contexts.

Note that the Interfaces page for ASA 5505 devices presents two tabbed panels: Hardware Ports and Interfaces. Similarly, the Interfaces page displayed for both FWSMs (version 3.1 and later) and ASAs (version 8.4.1 and later), operating in transparent mode also presents two tabbed panels: Interfaces and Bridge Groups.

Step 4 Add, edit and delete interfaces and related options, as necessary.

The Interfaces pages/panels and the Bridge Groups and Hardware Ports panels present standard Security Manager tables, with Add Row, Edit Row and Delete Row buttons, which are described in Using Tables, page 1-48.

The actual dialog box presented when you click the Add Row or Edit Row button depends on the type of device (and panel) you have selected. Refer to the following topics for device-specific dialog box information:

- Add/Edit Interface Dialog Box (PIX 6.3), page 46-28
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31
- Configuring Hardware Ports on an ASA 5505, page 46-61
- Add/Edit Bridge Group Dialog Box, page 46-62

Step 5 To manage Advanced Interface settings, including enabling communication between interfaces with the same security level, click the Advanced button at the bottom of the Interfaces page to open the Advanced Interface Settings dialog box. See Advanced Interface Settings (PIX/ASA/FWSM), page 46-68 for more information.

Step 6 When you are finished adding, editing and deleting interfaces, click **Save** at the bottom of the window to save your interface definitions to the Cisco Security Manager server.

Add/Edit Interface Dialog Box (PIX 6.3)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Table 46-7 Add/Edit Interface Dialog Box (PIX 6.3)

Element	Description
Enable Interface	Enables this interface to pass traffic. In addition to this setting, you must specify an IP Type and a Name before traffic can pass according to your security policy. You must enable a physical interface before traffic can pass through any enabled subinterfaces.
Type	Choose the type of interface: • **Physical** – VLAN is on the same network as its underlying hardware interface. • **Logical** – VLAN is associated with a logical interface.
Name	Provide an interface name up to 48 characters in length. The Name should be a memorable name for the interface that relates to its use. Supported interface names are: • Inside—Connects to your internal network. Must be most secure interface. • DMZ—Demilitarized zone (Intermediate interface). Also known as a perimeter network. • Outside—Connects to an external network or the Internet. Must be least secure interface.

Table 46-7 Add/Edit Interface Dialog Box (PIX 6.3) (continued)

Element	Description
Hardware Port	When defining a physical network interface, this value represents the name identifies the interface type and its slot or port in the device. When you add a logical network interface, you can choose any enabled physical interface to which you want to add a logical interface. If you do not see the desired hardware port, verify that the interface is enabled. Valid values are: • ethernet0 to ethernet*n*. • gb-ethernet*n*. where *n* represents the number of network interfaces in the device.
IP Type	The IP Type defines the type of IP addressing used for the interface; choose **Static IP** or **Use DHCP**, as described in Device Interface: IP Type (PIX 6.3), page 46-30. (The PPPoE option is not applicable to PIX 6.3 devices.) **Note** You can configure DHCP only on the outside interface of a security appliance.
Speed and Duplex	Lists the speed options for a physical interface; not applicable to logical interfaces. Choose one of the following options: • auto – Sets Ethernet speed automatically. The auto keyword can be used only with the Intel 10/100 automatic speed-sensing network interface card. • 10baset – 10-Mbps Ethernet half-duplex. • 10full – 10-Mbps Ethernet full-duplex. • 100basetx – 100-Mbps Ethernet half-duplex. • 100full – 100-Mbps Ethernet full-duplex. • 1000auto – 1000-Mbps Ethernet to auto-negotiate full- or half -duplex. **Tip** We recommend that you do not use this option to maintain compatibility with switches and other devices in your network. • 1000full – Auto-negotiate, advertising 1000-Mbps Ethernet full-duplex. • 1000full nonnegotiate – 1000-Mbps Ethernet full-duplex. • aui – 10-Mbps Ethernet half-duplex communication with an AUI cable interface. • bnc – 10-Mbps Ethernet half-duplex communication with a BNC cable interface. **Note** We recommend that you specify the speed of the network interfaces in case your network environment includes switches or other devices that do not handle autosensing correctly.

Table 46-7 Add/Edit Interface Dialog Box (PIX 6.3) (continued)

Element	Description
MTU	Specify the maximum packet size in bytes; that is, the maximum transmission unit (MTU). The value depends on the type of network connected to the interface. Valid values are 300 to 65535 bytes. Default is 1500.
Physical VLAN ID	For a physical interface, enter the VLAN ID, between 1 and 4094. This VLAN ID must not be in use on connected devices.
Logical VLAN ID	Provide the alias, a value between 1 and 4094, for the VLAN associated with this logical interface. This value is required if the logical interface Type is selected.
Security Level	Specify the security level of the interface: enter a value between 0 (least secure) and 100 (most secure). The security appliance lets traffic flow freely from an inside network to an outside network (lower security level). Many other security features are affected by the relative security level of two interfaces. • The *outside* interface is always 0. • The *inside* interface is always 100. • DMZ interfaces are between 1 and 99.
Roles	For more information on roles and how to define and use them, see Understanding Interface Role Objects, page 6-73. All interface roles assigned to this interface are listed in this field. Role assignments are based on pattern matching between the Name given to this interface and all currently defined Interface Role objects in Cisco Security Manager. Interface role objects are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. For more information on roles and how to define and use them, see Understanding Interface Role Objects, page 6-73.

Device Interface: IP Type (PIX 6.3)

A PIX 6.3 security device requires IP addressing for its interfaces; however, firewall interfaces do not have IP addresses until you assign them.

The Add/Edit Interface dialog box presented for a PIX 6.3 security device includes the section **IP Type**, where you specify the type of IP addressing for the interface and provide related parameters, as described here. See Add/Edit Interface Dialog Box (PIX 6.3), page 46-28 for information about the other sections of the dialog box.

Note The IP Type options presented for other security appliances are described in Device Interface: IP Type (PIX/ASA 7.0+), page 46-58.

Step 1 In the Add/Edit Interface dialog box, choose a method for address assignment from the **IP Type** list, and then provide related parameters, as follows:

- **Static IP** – Provide a static **IP Address** and **Subnet Mask** that represents the security device on this interface's connected network. The IP address must be unique for each interface.

  The Subnet mask can be expressed in dotted decimal format (for example, 255.255.255.0), or by entering the number of bits in the network mask (for example, 24). Beginning from Version 4.13, Cisco Security Manager allows 255.255.255.254 for a point to point interface. Do not use 255.255.255.255 for an interface connected to the network because this will stop traffic on that interface. If you omit the Subnet Mask value, a "classful" network is assumed, as follows:

  - The Class A netmask (255.0.0.0) is assumed if the first octet of the IP Address is 1 through 126 (i.e., addresses 1.0.0.0 through 126.255.255.255).
  - The Class B netmask (255.255.0.0) is assumed if the first octet of the IP Address is 128 through 191 (i.e., addresses 128.0.0.0 through 191.255.255.255).
  - The Class C netmask (255.255.255.0) is assumed if the first octet of the IP Address is 192 through 223 (i.e., addresses 192.0.0.0 through 223.255.255.255).

  Note Do not use addresses previously used for routers, hosts, or any other firewall device commands, such as an IP address in the global pool or a static NAT entry.

- **Use DHCP** – Enables Dynamic Host Configuration Protocol (DHCP) for automatic assignment of an IP address from a DHCP server on the connected network. The following options become available:
  - **Obtain Default Route using DHCP** – Check this box to obtain a default route from the DHCP server so that you do not need to configure a default static route.
  - **Retry Count** – The number of times the PIX will resend the DHCP request. Valid values are 4 to 16; the default is 2
- **PPPoE (PIX and ASA 7.2+)** – This option does not apply to PIX 6.3 devices.

Note You can configure DHCP only on the outside interface of a firewall device.

Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

These Add Interface and Edit Interface dialog boxes are used to define and configure interfaces, subinterfaces, redundant, and EtherChannel interfaces on PIX 7.0+, ASA, and FWSM devices. You can access the Add/Edit Interface dialog boxes from the Interfaces page. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.

Note The ASA 5505, combining switch and security appliance features, is a special case in that you configure both physical switch ports and logical VLAN interfaces. Thus, the Interfaces page displayed for ASA 5505 devices presents two tabbed panels: **Hardware Ports** and **Interfaces**. See Understanding ASA 5505 Ports and Interfaces, page 46-6 for more information.

ASA 8.4.1+ and FWSM 3.1+ devices operating in transparent mode also present two tabbed panels: **Interfaces** and **Bridge Groups**. See Add/Edit Bridge Group Dialog Box, page 46-62 for information about configuring bridge groups.

Many of the parameters presented in these dialog boxes vary according to device type and version, operational mode (routed versus transparent), and whether the device hosts a single or multiple contexts.

Note If you intend to use an interface for failover, you can define that interface in the Add Interface dialog box but do not configure it; instead, use the Failover page (see Chapter 50, "Configuring Failover"). In particular, do not specify an interface Name, as this parameter disqualifies the interface from being used as the failover link.

Using the Add Interface and Edit Interface Dialog Boxes

The following steps outline the general use of these dialog boxes:

1. An interface Type drop-down list appears at the top of the Add Interface and Edit Interface dialog boxes.

   Note Catalyst 6500 services modules (ASA-SMs and FWSMs) and the ASA 5505 do not present the Type list.

   Depending on device type, operating-system version and operating mode (router or transparent), the Type options presented will be two, three or all of the following:

   - **Physical Interface** – Choose this option to configure a physical interface on the device.
   - **Sub-Interface** – Choose this option to configure a logical interface (or VLAN connection) associated with a previously defined physical interface. Refer to Configuring Subinterfaces (PIX/ASA), page 46-7 for more information.
   - **Redundant** – Choose this option to configure two physical interfaces as a single logical "redundant interface." Refer to Configuring Redundant Interfaces, page 46-8 for more information.
   - **EtherChannel** – Choose this option to configure a logical interface consisting of a bundle of up to eight individual Ethernet links; this bundle is known as an EtherChannel, or a port-channel interface. (This option is available only on ASA 8.4+ devices.) Refer to Configuring EtherChannels, page 46-9, for more information.
   - **VNI Interface** – Choose this option to configure a VNI interface. They are virtual interfaces that keep network traffic separated on a given physical interface by using tagging. You apply your security policy directly to each VNI interface. All VNI interfaces are associated with the same VTEP interface. Refer to Configuring VNI Interfaces, page 46-15, for more information.
   - **Tunnel**– Choose this option to configure a logical interface -VTI, to support route based VPN method for the Site-to-Site VPN topologies. Refer to Configuring Tunnel Interface, page 46-22, for more information.

2. Below the Type option, the dialog boxes present up to three tabbed panels. Again, this depends on device type, operating-system version and operating mode.

   The PIX 7.0+ Add Interface and Edit Interface dialog boxes present two tabbed panels: General and Advanced. The ASA 7.0+ Add Interface and Edit Interface dialog boxes present three tabbed panels: General, Advanced and IPv6.

- Configure the General options, as appropriate. This panel is described in Add/Edit Interface Dialog Box: General Tab (PIX 7.0+/ASA/FWSM), page 46-33.
- Configure the Advanced-panel options, as appropriate. This panel is described in Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41.
- Configure the IPv6 options, as appropriate. This panel is described in Configuring IPv6 Interfaces (ASA/FWSM), page 46-47.

3. When you have finished configuring this interface, click **OK** to close the dialog box and return to the device Interfaces page.

Add/Edit Interface Dialog Box: General Tab (PIX 7.0+/ASA/FWSM)

The Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31, is used to define and configure interfaces, subinterfaces, VLAN interfaces, and redundant, and EtherChannel interfaces on firewall devices. You can access the Add/Edit Interface dialog box from the Interfaces page. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.

Note In the following descriptions, the term "interface" may be used generically to refer to any of these types of interface.

The General panel of this dialog box is used to configure general interface settings, including Name, Security Level and IP Type parameters. Note that many of the parameters presented in this panel vary according to device type and version, operational mode (routed versus transparent), and whether the device hosts a single or multiple contexts. Thus, some of the options in the following table may not appear for the device you are configuring.

Related Topics

- Configuring Subinterfaces (PIX/ASA), page 46-7
- Configuring Redundant Interfaces, page 46-8
- Configuring EtherChannels, page 46-9
- Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41
- Configuring IPv6 Interfaces (ASA/FWSM), page 46-47
- Understanding ASA 5505 Ports and Interfaces, page 46-6
- Configuring Hardware Ports on an ASA 5505, page 46-61

Table 46-8 General tab: Add/Edit Interface Dialog Box

Element	Description
Enable Interface	Enables this interface to pass traffic. By default, all physical interfaces are shut down. Traffic cannot traverse an interface of any type if the interface is not enabled. If you are defining a logical interface such as a subinterface, enable the physical interface it will be associated with before defining the subinterface. If you are defining a redundant interface or an EtherChannel interface, enable the member interfaces before defining the group interface. When you check this option, you must also specify a Name and, in routed mode, an IP Type (or IP Address and Subnet Mask on an FWSM or ASA-SM) before traffic can pass according to your security policy. In multiple-context mode, if you allocate a physical or logical interface to a context, the interface is enabled by default in the context. However, before traffic can pass through the context interface, you must also enable the interface in the system configuration. If you shut down an interface in the system execution space, that interface is shut down in all contexts in which it shared.
Management Only	Reserves this interface for device administration. Only traffic for management of this device is accepted; pass-through traffic for other interfaces and devices is rejected. You cannot set a Primary or Secondary ISP interface to be Management Only. Defining a management-only EtherChannel interface has certain member-interface restrictions. See Configuring EtherChannels, page 46-9, for more information. **Note** This is not available on devices in transparent mode. If an interface is assigned as Management Only, then Route Map cannot be assigned to that interface. In other words, either Management Only or Route Map can be assigned to an interface but not both.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
Interface	On the ASA 5505, the Hardware Port is specified on the Hardware Ports panel (see Configuring Hardware Ports on an ASA 5505, page 46-61). Also, this option is not part of Catalyst 6500 services module (ASA-SM and FWSM) configuration. For a physical interface, provide the specific hardware port assigned to the interface: enter a physical port ID, which includes network type, slot and port number, in the form: *type[slot/]port*. This is also the name by which subinterfaces can be associated with the interface. The network type specified for the physical interface can be either Ethernet or GigabitEthernet; on the ASA 5580, TenGigabitEthernet is also available. This field provides automatic pattern matching: if you begin typing with the letter e, "Ethernet" is inserted into the field. Similarly, typing the letter g produces "GigabitEthernet." Therefore, valid values are: • Ethernet0 to Ethernet*n* • GigabitEthernet0 to GigabitEthernet*n* • GigabitEthernet*s*/*n* • TenGigabitEthernet*s*/*n* (ASA 5580 only) where *s* represents a slot number, and *n* represents a port number, up to the maximum number of network ports in the slot or device. For an ASA 5500 series appliance, enter the type and a slot/port pair; for example, *gigabitethernet0/1*. Ports that are built into the chassis are assigned to slot 0, while ports on the 4-Port Gigabit Ethernet Security Services Module (4GE SSM) are assigned to slot 1. When you enter a slot/port pair, the Media Type options are enabled. The ASA 5500 series appliances also include a **management** interface type. The management interface is a Fast Ethernet interface designed for device-management traffic only, and is specified as *management0/0*. You can, however, use this physical interface for through traffic if desired (be sure the Management Only option is not selected). Thus, in transparent firewall mode, you can use the management interface in addition to the two interfaces allowed for through traffic. You can also add subinterfaces to the management interface to provide management in each security context in multiple-context mode. If you are defining a subinterface, you can simply choose the desired Hardware Port from a list of previously defined ports (you must also supply a VLAN ID). If you do not see a desired interface ID, be sure that Interface is defined and enabled.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
Name	Provide an identifier for this interface of up to 48 characters in length. The name should be a memorable name for the interface that relates to its use. However, if you are using failover, do not name interfaces that you are reserving for failover communications; this includes an EtherChannel intended for failover, as well as its member interfaces. Also, do not name interfaces intended for use as a member of a redundant-interface pair. Certain names are reserved for specific interfaces, in accordance with the interface naming conventions of the security appliance. As such, these reserved names enforce default, reserved security levels, as follows: • Inside – Connects to your internal network. Must be the most secure interface. • DMZ – "Demilitarized zone" attached to an intermediate interface. DMZ is also known as a perimeter network. You can name a DMZ interface any name you choose. Typically, DMZ interfaces are prefixed with "*DMZ*" to identify the interface type. • Outside – Connects to an external network or the Internet. Must be the least secure interface. Similarly, a subinterface name typically identifies its associated interface, in addition to its own unique identifier. For example, *DMZoobmgmt* could represent an out-of-band management network attached to the DMZ interface. **Note** Again, **do not name the interface if you intend to use it for failover, or as a member of a redundant interface**. See Chapter 50, "Configuring Failover" and Configuring Redundant Interfaces, page 46-8 for more information.
Security Level	Specify the security level of the interface: enter a value between 0 (least secure) and 100 (most secure). The security appliance lets traffic flow freely from an inside network to an outside network (lower security level). Many other security features are affected by the relative security level of two interfaces. • The *outside* interface is always 0. • The *inside* interface is always 100. • DMZ interfaces are between 1 and 99.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
Media Type	When Interface is the chosen Type and you enter a hardware port ID with slot/port numbers in the Hardware Port field, these options are enabled. (These options apply to ASA slot/port interfaces only.) For all ASA 5500 series appliances, except the 5505, ports that are built into the chassis are assigned to slot 0, while ports on the 4GE SSM are assigned to slot 1. By default, all connectors used on an ASA are RJ-45 connectors. However, the ports on the 4GE SSM can include fiber SFP connectors. As part of the interface configuration for one of these fiber-based connections, you must change the Media Type setting from the default (RJ45) to the fiber-connector setting (SFP). Fiber-based interfaces do not support duplexing and have a fixed speed, so the Duplex option is disabled, and the Speed options are limited to auto and nonegotiate. Select the connector type used by this slot-1 interface: • **RJ45** – The port uses RJ-45 (copper) connectors. • **SFP** – The port uses fiber SFP connectors. Required for 10-Gigabit Ethernet cards.
VLAN ID	When Subinterface is the chosen interface Type, or when you are defining a logical interface on a device operating in transparent mode, on an ASA 5505, or on a Catalyst 6500 services module (ASA-SM or FWSM), provide a VLAN ID for this interface. For PIX/ASA devices running operating system 7.2(2)18 or earlier, valid VLAN IDs are 1 to 1001; with version 7.2(2)19 or later, valid IDs are 1 to 4090. For Catalyst 6500 services modules, valid IDs are 1 to 4096. The specified VLAN ID must not be in use on any connected device. Some VLAN IDs might be reserved on connected switches; see the switch documentation for more information. In multiple-context mode, you can only set the VLAN ID in the system configuration. See Configuring Subinterfaces (PIX/ASA), page 46-7 for more information.
Subinterface ID	When Subinterface is the chosen interface Type, or when defining an interface on a device operating in transparent mode, provide an integer between 1 and 4294967293 as the Subinterface ID. For subinterface port identification, this ID is appended to the chosen Hardware Port. For example, *GigabitEthernet0.4* represents the subinterface assigned an ID of 4, operating on the port GigabitEthernet0. **Note** You cannot change the Subinterface ID after you set it.
Route Map	Select the Route Map from the Route Map Object Selector dialog box. **Note** Except VNI Interface, all other interface types support Policy Based Routing for ASA devices running the software version 9.4(1) or higher. VNI Interface supports Policy Based Routing for ASA devices running the software version 9.5(1) or higher.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
IP Type	**PIX 7.0+ and ASA (except the 5505 in transparent mode) only.** The IP Type defines the type of IP addressing used for the interface; choose **Static IP**, **Use DHCP**, or **PPPoE** (as described in Device Interface: IP Type (PIX/ASA 7.0+), page 46-58). **Note** You can configure DHCP and PPPoE only on the outside interface of a security appliance.
IP Address Subnet Mask	**Catalyst 6500 services modules (ASA-SMs and FWSMs) in routed mode only.** Use these two fields to assign an IP address and subnet mask to the VLAN interface. The IP address must be unique for each interface. The Subnet Mask can be expressed in dotted decimal format (for example, 255.255.255.0), or by entering the number of bits in the network mask (for example, 24). Till Version 4.12, 255.255.255.254 and 255.255.255.255 were not to be used for an interface connected to the network because it would stop traffic on that interface. Beginning from Version 4.13, /31 subnet mask (or 2555.2555.255.254) is supported for a point to point interface connected to the network. Cisco Security Manager displays a warning message on saving the interface record. If you omit the Subnet Mask value, a "classful" network is assumed, as follows: • The Class A netmask (255.0.0.0) is assumed if the first octet of the IP Address is 1 through 126 (that is, addresses 1.0.0.0 through 126.255.255.255). • Subnet Mask The Class B netmask (255.255.0.0) is assumed if the first octet of the IP Address is 128 through 191 (that is, addresses 128.0.0.0 through 191.255.255.255). • The Class C netmask (255.255.255.0) is assumed if the first octet of the IP Address is 192 through 223 (that is, addresses 192.0.0.0 through 223.255.255.255). **Note** Do not use addresses previously used for routers, hosts, or any other firewall device commands, such as an IP address in the global pool or a static NAT entry.
Description	You can enter an optional description of up to 240 characters on a single line, without carriage returns. In multiple-context mode, the system description is independent of the context description. For a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
Redundant Interface; these options not available on ASA 5505 devices or Catalyst 6500 services modules (ASA-SMs and FWSMs).	
Redundant ID	When Redundant Interface is the chosen interface Type, provide an identifier for this redundant interface; valid IDs are the integers from 1 to 8. See Configuring Redundant Interfaces, page 46-8 for more information.
Primary Interface Secondary Interface	When Redundant Interface is the chosen interface Type, choose the primary member of the redundant interface pair from the Primary Interface list of available interfaces. Available interfaces are presented by Hardware Port IDs, as named interfaces cannot be used for a redundant interface pair. Similarly, choose the secondary member of the redundant interface pair from the Secondary Interface list of available interfaces. **Note** Member interfaces must be enabled and of the same type (e.g., GigabitEthernet), and cannot have a Name, IP Address, or Security Level assigned. In fact, do not configure any options other than Duplex and Speed on the member interfaces.
These options available on ASA 5505 devices only.	
Block Traffic To	Restricts this VLAN interface from initiating contact with the VLAN chosen here.
Backup Interface	Choose a VLAN interface as a backup interface, for example, to an ISP. The backup interface does not pass traffic unless the default route through the primary interface fails. To ensure that traffic can pass over the backup interface, be sure to configure default routes on both the primary and backup interfaces so that the backup interface can be used when the primary fails.
Active MAC Address Standby MAC Address	Use the Active MAC Address field to manually assign a private MAC address to the interface; the Standby MAC Address field can be used to set a standby MAC address for use with device-level failover. Refer to Device Interface: MAC Address, page 46-60 for more information about these fields.
EtherChannel Interface options; available on ASA 8.4.1+ devices only.	
EtherChannel ID	When EtherChannel is the chosen interface Type, enter an identifier for this EtherChannel (also referred to as a “port-channel”). Valid values are 1 to 48—you can define up to 48 port-channel groups. See Configuring EtherChannels, page 46-9, for more information.

Table 46-8 General tab: Add/Edit Interface Dialog Box (continued)

Element	Description
Available Interfaces/Members in Group	When EtherChannel is the chosen interface Type, you can assign interfaces to this EtherChannel group by selecting them in the Available Interfaces list and then clicking the >> button to add them to the members list to the right. You can assign up to 16 interfaces to a channel group. For ASA 9.2(1) and later, each channel group can have up to 16 active interfaces. For switches that support only eight active interfaces and for ASA versions earlier than 9.2(1), only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure. Alternatively, you can create a static EtherChannel by setting LACP Mode to On (on the Advanced tab, see Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41), which means all interfaces in the group can pass traffic. **Note** All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the type and speed for the group. See Configuring EtherChannels, page 46-9, for more information.

Add/Edit Interface Dialog Box: Cisco Firepower 9000 (General and Advanced tabs)

For the elements supported in Cisco Firepower 9000 devices for the General and Advanced tabs, see the Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31. In addition, the following changes are applicable only to Cisco Firepower 9000 devices.

Table 46-9 Add/Edit Interface Dialog Box Cisco Firepower 9000

Element	Description
Type	Choose the type of interface. Redundant Interfaces are not supported in Cisco Firepower 9000 devices.
Management Only Individual	Applicable only in Cisco Firepower 9000 devices and only if the device is in cluster mode. **Note** You cannot enable both Management Only and Management Only Individual check boxes at the same time. You can configure Cluster pool only when Management Only Individual check box is selected.

Element	Description
Name	Provide an interface name up to 48 characters in length. The Name should be a memorable name for the interface that relates to its use. The Interface name must begin with “Ethernet” and must have the following format: Ethernet[slot]/[port]/sub-port, where, • slot is between 1 and 3 • port is between 1 and 8 • sub-port is between 1 and 4 • sub-port is not applicable for slot 1
The following elements are not supported in Cisco Firepower 9000 devices:	
Media Type (General Tab)	
Duplex(Advanced Tab)	
Speed (Advanced Tab)	
Available Interfaces/Members in Group (General Tab)	
Load Balancing (Advanced Tab)	
LACP Mode (Advanced Tab)	
VSS or vPC Switch ID (Advanced Tab)	
Active Physical Interfaces (Advanced Tab)	
Span EtherChannel across the ASA Cluster (Advanced Tab)	
Enable load balancing between switch pairs in VSS or vPC mode (Advanced Tab)	
Member Interface Configuration (Advanced Tab)	

Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+)

The Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31, is used to define and configure interfaces, subinterfaces, redundant, and EtherChannel interfaces on ASA and PIX 7.0+ devices. You can access the Add/Edit Interface dialog box from the Interfaces page. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.

The Advanced panel of this dialog box is used to configure basic interface settings, including Duplex, Speed, and maximum transmission unit (MTU) parameters, as described in the following table.

Related Topics

- Add/Edit Interface Dialog Box: General Tab (PIX 7.0+/ASA/FWSM), page 46-33
- Configuring IPv6 Interfaces (ASA/FWSM), page 46-47

Table 46-10 Advanced tab: Add/Edit Interface Dialog Box (ASA/PIX 7.0+)

Element	Description
Duplex	Lists the duplex options for the interface, including Auto, Full, Half, or N/A, depending on the interface type. For TenGigabitEthernet (ASA 5580 only), Duplex is automatically set to Full. **Note** This option is not available when Subinterface or Redundant is the chosen Interface type.
Speed	Lists the speed options (in bits per second) for a physical interface; not applicable to logical interfaces. The speeds available depend on the interface type. • auto • 10 • 100 • 1000 • 10000 (set automatically for a TenGigabitEthernet interface; available only on ASA 5580) • nonegotiate **Note** This option is not available when Subinterface or Redundant is the chosen Interface type.
MTU	Specify the maximum packet size in bytes; that is, the maximum transmission unit (MTU). The value depends on the type of network connected to the interface. Valid values are 300 to 65535 bytes. Default is 1500 for all types except PPPoE, for which the default is 1492. In multiple-context mode, set the MTU in the context configuration.
Active MAC Address Standby MAC Address	**Available only on PIX 7.2+ and ASA 7.2+ devices.** Use the Active MAC Address field to manually assign a private MAC address to the interface; the Standby MAC Address field can be used to set a standby MAC address for use with device-level failover. Refer to Device Interface: MAC Address, page 46-60 for more information about these fields.
Roles	All interface roles assigned to this interface are listed in this field. Role assignments are based on pattern matching between the Name given to this interface and all currently defined Interface Role objects in Cisco Security Manager. Interface role objects are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. For more information on roles and how to define and use them, see Understanding Interface Role Objects, page 6-73.
MAC Address	Site specific MAC address.
Site ID	Site ID to specify the site the current unit belongs to.

Table 46-10 Advanced tab: Add/Edit Interface Dialog Box (ASA/PIX 7.0+) (continued)

Element	Description
Beginning with Security Manager version 4.9 for ASA devices running the software version 9.5(1) or higher, you can use inter-site clustering for Spanned EtherChannels in routed mode. To avoid MAC address flapping, configure a site ID for each cluster member so that a site-specific MAC address for each interface can be shared among a site's units.	
EtherChannel Interface options; available on ASA 8.4.1+ devices only.	
Load Balancing	When EtherChannel is the chosen interface Type (on the General panel), choose a load-balancing method for the channel links. See About EtherChannel Load Balancing, page 46-13, for more information about this option.
LACP Mode	Select the desired **LACP Mode**; the default is Active, which means up to eight interfaces are active, while up to eight are in stand-by mode, as determined by the Minimum and Maximum values under Active Physical Interfaces. If you select On, a static port-channel is created in which all member interfaces are all "on," meaning you can have up to 16 ports passing traffic, with no stand-by ports. When you select this option, the Mode for all interfaces assigned to this EtherChannel group is switched to On (if the Mode for each is not already On). See Editing LACP Parameters for an Interface Assigned to an EtherChannel, page 46-12, for more information about this mode.
Active Physical Interfaces	When EtherChannel is the chosen interface Type (on the General panel), specify the minimum and maximum number of interfaces that can be active for this EtherChannel group: • **Minimum** – Specify the minimum number of active interfaces for this group. For ASA 9.2(1)+, you can specify a value from 1 to 16; for earlier versions, enter a value from 1 to 8. If the active interfaces in the channel group falls below this value, then the port-channel interface goes down, and could trigger a device-level failover. • **Maximum** – Specify the maximum number of interfaces that can be active. For ASA 9.2(1)+, you can specify a value from 1 to 16; for earlier versions, enter a value from 1 to 8. For 16 active interfaces, be sure that your switch supports the feature (for example, the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module). If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer. Interfaces available to the channel are selected on the General tab of this dialog box (Add/Edit Interface Dialog Box: General Tab (PIX 7.0+/ASA/FWSM), page 46-33). Specifying 3, 5, 6, or 7 active ports in an EtherChannel bundle provides poor load balancing, because some ports get up to twice the load of others. We recommend specifying 2, 4, or 8 active ports per EtherChannel to achieve effective load balancing. (A value of 1 provides no load balancing at all.)

Table 46-10 Advanced tab: Add/Edit Interface Dialog Box (ASA/PIX 7.0+) (continued)

Element	Description
DHCP Relay options; available on ASA-SM 9.1.2+ devices only.	
DHCP Relay Servers	Enter the IP address or select a Networks/Hosts object representing the interface-specific DHCP server to which DHCP requests on this interface are relayed. Use a comma to separate multiple values. You can configure a maximum of 4 interface-specific DHCP relay servers and a maximum of 10 global and interface-specific DHCP relay servers combined. **Note** IPv6 is not supported for interface-specific servers. When a DHCP request enters an interface, the DHCP servers to which the ASA relays the request depends on your configuration. You can configure the following types of servers: • Interface-specific DHCP servers—When a DHCP request enters a particular interface, then the ASA relays the request only to the interface-specific servers. • Global DHCP servers—When a DHCP request enters an interface that does not have interface-specific servers configured, the ASA relays the request to all global servers. If the interface has interface-specific servers, then the global servers are not used. For more information, see DHCP Relay Page, page 52-5.
DHCP Relay Trust Info (Option 82)	Specifies that you want to trust this DHCP client interface. You can configure interfaces as trusted interfaces to preserve DHCP Option 82. **Note** You can also trust all DHCP client interfaces. For more information, see DHCP Relay Page, page 52-5. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface.

Secure Group Tagging options; available on ASA 9.3.1+ devices only.

SGT plus Ethernet Tagging, also called Layer 2 SGT Imposition, enables the ASA to send and receive security group tags on Ethernet interfaces using Cisco proprietary Ethernet framing (EtherType 0x8909), which allows the insertion of source security group tags into plain-text Ethernet frames. The ASA inserts security group tags on the outgoing packet and processes security group tags on the incoming packet, based on a manual per-interface configuration. This feature allows inline hop-by-hop propagation of endpoint identity across network devices and provides seamless Layer 2 SGT Imposition between each hop.

Note Supported only on physical interfaces, VLAN interfaces, port channel interfaces, and redundant interfaces. Not supported on logical interfaces or virtual interfaces, such as BVI, TVI, and VNI.

Does not support failover links or cluster control links.

Table 46-10 Advanced tab: Add/Edit Interface Dialog Box (ASA/PIX 7.0+) (continued)

Element	Description
Enable secure group tagging for Cisco TrustSec	Enables SGT plus Ethernet Tagging (also called Layer 2 SGT Imposition).
Tag egress packets with secure group tags	Enables propagation of a security group tag (called sgt) on an interface.
Assign a static secure group tag to all ingress packets	Applies a static security group tag to incoming traffic from the peer. If enabled, you must specify the SGT number to use in the Secure Group Tag field.
Secure Group Tag	Specifies the SGT number to apply to incoming traffic from the peer. Valid values are from 2-65519.
Trusted Interface	Indicates that ingress traffic on the interface should not have its existing SGT overwritten with the static SGT specified.
ASA Cluster (Layer 3); available on ASA 5580 and 5585 devices in cluster mode only. Supported by all interfaces when ASA cluster is in Router mode and supported by management interface when ASA cluster is in Transparent mode.	
IPv4 Address Pool	Enter or select the IPv4 Pool object that represents the pool of addresses to use.
MAC Address Pool	Enter or select the MAC Pool object that represents the pool of MAC addresses to use.
ASA Cluster (Layer 2); available on ASA 5580 and 5585 devices in cluster mode only. Supported on EtherChannel interfaces for ASA clusters. Not supported on Management interface when ASA cluster is in Transparent mode.	
Span EtherChannel across the ASA Cluster	Select to configure an EtherChannel that spans all ASAs in the cluster, and provides load balancing as part of the EtherChannel operation.
Enable load balancing between switch pairs in VSS or vPC mode	(Optional) If you are connecting the ASA to two switches in a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you should enable load balancing by checking the **Enable load balancing between switch pairs in VSS or vPC mode** check box. This feature ensures that the physical link connections between the ASAs to the VSS (or vPC) pair are balanced.
Member Interface Configuration	Identifies the LACP mode for the interface and the Virtual Switching System (VSS) or Virtual Port Channel (vPC) switch to which a given interface is connected, 1 or 2.
Advanced tab options specific to ASA 5505 devices (routed mode only)	
Block Traffic To	Restricts this VLAN interface from initiating contact with the VLAN chosen here.
Backup Interface	Choose a VLAN interface as a backup interface, for example, to an ISP. The backup interface does not pass traffic unless the default route through the primary interface fails. To ensure that traffic can pass over the backup interface, be sure to configure default routes on both the primary and backup interfaces so that the backup interface can be used when the primary fails.

Table 46-10 Advanced tab: Add/Edit Interface Dialog Box (ASA/PIX 7.0+) (continued)

Element	Description
Advanced tab options specific to FWSM 3.1+ devices	
Bridge Group	For an FWSM 3.1+ operating in transparent mode, this read-only field indicates the Bridge group to which this interface is assigned. See Add/Edit Bridge Group Dialog Box, page 46-62 for more information.
ASR Group	To add this interface to an asymmetric routing group, enter the ASR group number in this field. Stateful failover must be enabled for asymmetric routing support to function properly between units in failover configurations. Valid values for ASR group range from 1 to 32. See About Asymmetric Routing Groups, page 46-6 for more information.

Pause Frame for Flow Control options

When a network interface gets over loaded, flow control allows it to send PAUSE requests to the devices sending it data to allow the over loaded condition to clear. If flow control is not enabled and an over loaded condition occurs, the device will drop packets.

When the receiving part of the interface reaches the high water mark, the transmitting part of the interface starts to generate pause frames. The remote device is expected to stop / reduce the transmission of packets for the pause time mentioned in the pause frame. If the receiving part of the interface is able to clear its queue or reaches the low water mark within the pause time, the transmitting part of the interface sends out a special pause frame that mentions the pause time as zero.This enables the remote device to start to transmit packets. If the receiving part of the interface still works on the queue, once the pause time expires, the transmitting part of the interface sends a new pause frame again with a new pause time.

Note Pause Frame for flow control is supported only on physical interfaces on ASA 8.2 and above, in the single and multi-context mode. It is not supported on logical interfaces or virtual interfaces, such as BVI, TVI, and VNI.

Enable Pause Frame	(Optional) Enables transmission of pause frame for flow control.
Use Default Values	(Optional) Uses default values for Low Watermark, High Watermark and Pause Time, based on the device. If this is unchecked, specify the values as per the Device specific Pause Frame Flow Control values reference table.
Low Watermark (in Kilobytes)	Enter a value for the low-water mark. After the interface sends a pause frame, when the buffer usage is reduced below the low-water mark, the interface sends an “transmission on’ frame. The remote device can resume transmitting data.
High Watermark (in Kilobytes)	Enter a value for the high-water mark.When the buffer usage exceeds the high-water mark, the interface sends a pause frame.
Pause Time	Enter a value for the pause refresh threshold value, between 0 and 65535 slots. Each slot is the amount of time to transmit 64 bytes, so the time per unit depends on your link speed. The remote device can resume traffic after receiving an transmission on frame, or after the transmission off frame expires, as controlled by this timer value in the pause frame. If the buffer usage is consistently above the high-water mark, pause frames are sent repeatedly, controlled by the pause refresh threshold value.

Table 46-11 Device specific Pause Frame Flow Control values

Device Type	Low Watermark Range (in Kb)	Default Low Watermark (in Kb)	High Watermark Range (in Kb)	Deafult High Watermark (in Kb)	Range of Pause Time	Default Pause Time
ASA 5515	0-20	8	0-20	16	0-65535	26624
ASA 5525	0-20	8	0-20	16	0-65535	26624
ASA 5545	0-20	8	0-20	16	0-65535	26624
ASA 5510	0-48	16	0-48	24	0-65535	26624
ASA 5585	Values are not supported; only “flowcontrol send on” is supported.					
ASA 5506	1-25	3	1-25	8	1-65535	18432
ISA-3000-2C2F	0-64	27	0-64	34	0-65535	26624
ISA-3000-4C	0-64	27	0-64	34	0-65535	26624
1783-SAD4T0S	0-64	27	0-64	34	0-65535	26624

Configuring IPv6 Interfaces (ASA/FWSM)

When Interface, Subinterface, Redundant, or EtherChannel is the chosen Type in the Add Interface or Edit Interface dialog box, the dialog box presents three tabbed panels of options: General, Advanced and IPv6. The options provided by the **IPv6** panel are described in this section.

Note These options are available only on ASA 7.0+ devices in routed mode; ASA 8.2+ devices in transparent mode; and FWSM 3.1+ devices in routed mode.

Navigation Path

You can access the IPv6 panel in the Add Interface and Edit Interface dialog boxes, which are accessed from the ASA or FWSM Interfaces page, as described in Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26.

Related Topics

- IPv6 Support in Security Manager, page 1-8
- Add/Edit Interface Dialog Box: General Tab (PIX 7.0+/ASA/FWSM), page 46-33
- Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41

Field Reference

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM)

Element	Description
Enable IPv6	Check this box to enable IPv6 and configure IPv6 addresses on this interface. You can deselect this option to disable IPv6 on the interface, but retain the configuration information.

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM) (continued)

Element	Description
Enforce EUI-64	When selected, use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link is enforced. When this option is enabled on an interface, the source addresses of IPv6 packets received on the interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the interface identifier in an IPv6 packet is not in the Modified EUI-64 format, the packet is dropped and the following system log message is generated: `%PIX|ASA-3-325003: EUI-64 source address check failed.` Address format verification is performed only when a flow is created. Packets from an existing flow are not checked. Additionally, address verification can be performed only for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address. The Modified EUI-64 format interface identifier is derived from the 48-bit link-layer (MAC) address by inserting the hex number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link-layer address. To ensure the chosen address is from a unique Ethernet MAC address, the next-to-lowest order bit in the high-order byte is inverted (universal/local bit) to indicate the uniqueness of the 48-bit address. For example, an interface with a MAC address of 00E0.B601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM) (continued)

Element	Description
DAD Attempts	To specify the number of consecutive neighbor solicitation messages that are sent on an interface during duplicate address detection (DAD), enter a number from 0 to 600 in this field. Entering 0 disables duplicate address detection on the interface. Entering 1 configures a single transmission without follow-up transmissions; this is the default. Duplicate address detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection uses neighbor solicitation messages to verify the uniqueness of unicast IPv6 addresses. When duplicate address detection identifies a duplicate address, the state of the address is set to DUPLICATE and the address is not used. If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface and an error message similar to the following is issued: `%PIX-4-DUPLICATE: Duplicate address FE80::1 on outside` If the duplicate address is a global address of the interface, the address is not used and an error message is issued, similar to that shown previously for a duplicate link-local address. All configuration commands associated with the duplicate address remain as-configured while the state of the address is set to DUPLICATE. If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address, and all other IPv6 address associated with the interface are regenerated (that is, duplicate address detection is performed only on the new link-local address).
NS Interval	The interval between IPv6 neighbor solicitation retransmissions, in milliseconds. Valid values range from 1000 to 3600000 milliseconds; the default value is 1000 milliseconds. **Note** This value is included in all IPv6 router advertisements sent out on this interface.
Reachable Time	The amount of time, in milliseconds, within which a remote IPv6 node is considered still reachable, after initial reachability was confirmed. Valid values range from 0 to 3600000 milliseconds, the default value is 0. When 0 is used for the value, the reachable time is set as undetermined—it is up to the receiving devices to set and track reachable time. A configured time enables detection of unavailable neighbors. A shorter time allows detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.
Managed Config Flag	Whether or not to set the flag "managed-config-flag" in the IPv6 router advertisement packet.

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM) (continued)

Element	Description
Other Config Flag	Whether or not to set the flag "other-config-flag" in the IPv6 router advertisement packet.
Enable RA	When checked, IPv6 router advertisement transmissions are enabled on the interface. The following options are enabled: • **RA Lifetime** – The "router lifetime" value specifies how long nodes on the local link should consider the security appliance as the default router on the link. Valid values range from 0 to 9000 seconds; the default is 1800 seconds. Entering 0 indicates that the security appliance should not be considered a default router on the selected interface. Any non-zero value should not be less than the following RA Interval value. **Note** This value is included in all IPv6 router advertisements sent out on this interface. • **RA Interval** – The interval between IPv6 router advertisement transmissions on this interface. Valid values range from 3 to 1800 seconds, (or from 500 to 1800000 milliseconds if the following RA Interval in Milliseconds option is checked); the default is 200 seconds. The interval between transmissions should be less than or equal to the RA Lifetime value if it is non-zero. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value. • **RA Interval in Milliseconds** – Checking this option indicates that the provided RA Interval value is in milliseconds, rather than seconds.

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM) (continued)

Element	Description
Interface IPv6 Addresses	The IPv6 addresses assigned to the interface are specified in this section of the dialog box. • **Link-Local Address** – To override the link-local address that is automatically generated for the interface, enter the desired IPv6 link-local address in this field. The link-local address is composed of the link-local prefix FE80::/64 and the interface ID in Modified EUI-64 format. For example, an interface with a MAC address of 00E0.B601.3B7A would have a link-local address of FE80::2E0:B6FF:FE01:3B7A. An error will occur if another host is using the specified address. • **Enable Address Auto-Configuration** – Select this option to enable automatic configuration of IPv6 addresses on the interface using stateless autoconfiguration. The addresses are configured based on the prefixes received in Router Advertisement (RA) messages. If a link-local address has not been configured, then one is automatically generated for this interface. An error occurs if another host is already using the generated link-local address. – **Trust the DHCP Servers for default gateway**– Select this radio button to install a default route from Router Advertisements that come from a trusted source - the directly-connected network. – **Ignore trust and accept router advertisements** – Select this radio button to install a default route from Router Advertisements that come from another network. • The table in this section displays the IPv6 addresses assigned to this interface. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Address for Interface Dialog Box, page 46-52.
Interface IPv6 Prefixes	Use the table in this section to configure which IPv6 prefixes (that is, the network portion of the IPv6 addresses) are included in IPv6 router advertisements. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Prefix Editor Dialog Box, page 46-54.

Table 46-12 IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM) (continued)

Element	Description
Interface IPv6 DHCP	Use this section to enable the DHCPv6 Prefix Delegation client on one or more interfaces. The ASA obtains one or more IPv6 prefixes that it can subnet and assign to inside networks. Typically, the interface on which you enable the prefix delegation client obtains its IP address using the DHCPv6 address client; only other ASA interfaces use addresses derived from the delegated prefix. Select one of the following: • **Server Pool** – Select this to configure the IPv6 DHCP pool that contains the information you want the DHCPv6 server to provide. You can configure separate pools for each interface if you want, or you can use the same pool on multiple interfaces. Use the Add Row and Edit Row buttons in the **DHCP Pool Selector** dialog to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the Add or Edit DHCPv6 Pool Dialog Box, page 46-56. OR • **Client Prefix Delegation Name** – Enable the DHCPv6 Prefix Delegation client by entering a name for the prefix(es) obtained on this interface. Valid values are a string not exceeding 200 characters. – **DHCPv6 Prefix Hint** – Use the Add Row button to provide one or more hints about the delegated prefix that you want to receive. Typically you want to request a particular prefix length, such as ::/60, or if you have received a particular prefix before and want to ensure you get it again when the lease expires, you can enter the whole prefix as the hint. If you enter multiple hints (different prefixes or lengths), then it is up to the DHCP server which hint to honor, or whether to honor the hint at all. **Note** If the prefix suggested as the hint is a valid prefix in the associated local prefix pool and is not assigned elsewhere, the server delegates the client-suggested prefix. Otherwise, the hint is ignored and a prefix is delegated from the free list in the pool. – **Enable DHCP** – Select this to obtain an address using DHCPv6. Optionally, select **Enable Default Route** to obtain a default route from Router Advertisements.

Note If a DHCPv6 client or Server Pool is configured on an IPv6 Interface, the same interface cannot be used to configure DHCPv6 Relay.

IPv6 Address for Interface Dialog Box

This dialog box is used to add or edit an IPv6 address assigned to an ASA or FWSM interface. Multiple IPv6 addresses can be assigned to the interface in the IPv6 panel of the Add Interface or Edit Interface dialog box.

Note This dialog box is available only on ASA 7.0+ devices in routed mode; ASA 8.2+ devices in transparent mode; and FWSM 3.1+ devices in routed mode.

Navigation Path

You can access the IPv6 Address for Interface dialog box:

- From the IPv6 panel of the ASA or FWSM Add Interface and Edit Interface dialog boxes.
- From the Management IPv6 page of an ASA 5505 in transparent firewall mode (version 8.2 and 8.3 devices only).

Click the Add Row or Edit Row buttons beneath the table in the Interfaces IPv6 Addresses section to open the dialog box.

Related Topics

- IPv6 Prefix Editor Dialog Box, page 46-54
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26
- Management IPv6 Page (ASA 5505), page 47-11

Field Reference

Table 46-13 IPv6 Address for Interface Dialog Box

Element	Description
Prefix Name	(Optional) Enter a prefix name to use a delegated prefix. Valid values are a string not exceeding 200 characters. **Tip** 'DHCP' is a reserved word; Cisco Security Manager will not accept it as Prefix Name. **Note** Make sure that the DHCPv6 Prefix Delegation Client is enabled on this ASA Interface. For more information, see the Interface IPv6 DHCP element in the Table 46-12IPv6 tab: Add/Edit Interface Dialog Box (ASA/FWSM), page 46-47>

Table 46-13 IPv6 Address for Interface Dialog Box (continued)

Element	Description
Address/Prefix Length	Enter an IPv6 network address to be assigned to the interface, with its Prefix Length appended, where the Prefix Length integer indicates how many of the high-order, contiguous bits of the address comprise network portion of the address. A slash (/) must precede the Prefix Length. For example, 3FFE:C00:0:1::/64. Typically, the delegated prefix will be /60 or smaller so you can subnet to multiple /64 networks. /64 is the supported subnet length if you want to support SLAAC for connected clients. You should specify an address that completes the /60 subnet, for example ::1:0:0:0:1. Enter :: before the address in case the prefix is smaller than /60. For example, if the delegated prefix is 2001:DB8:1234:5670::/60, then the global IP address assigned to this interface is 2001:DB8:1234:5671::1/64. The prefix that is advertised in router advertisements is 2001:DB8:1234:5671::/64. In this example, if the prefix is smaller than /60, the remaining bits of the prefix will be 0's as indicated by the leading ::. For example, if the prefix is 2001:DB8:1234::/48, then the IPv6 address will be 2001:DB8:1234::1:0:0:0:1/64.
EUI-64	If this box is checked, the EUI-64 interface ID will be used in the low-order 64 bits of the IPv6 address. If the value specified for the Prefix Length is greater than 64 bits, the prefix bits have precedence over the interface ID. An error occurs if another host is using the specified address. The Modified EUI-64 format interface ID is derived from the 48-bit link-layer (MAC) address by inserting the hex number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link layer address. To ensure the chosen address is from a unique Ethernet MAC address, the next-to-lowest order bit in the high-order byte is inverted (universal/local bit) to indicate the uniqueness of the 48-bit address. For example, an interface with a MAC address of 00E0.B601.3B7A would have a 64 bit interface ID of 02E0:B6FF:FE01:3B7A.
IPv6 Address Pool	Enter or select the IPv6 Pool object that represents the pool of addresses to use.

IPv6 Prefix Editor Dialog Box

This dialog box is used to add or edit an IPv6 prefix (that is, the network portion of an IPv6 address), providing control over individual parameters, including whether the prefix should be included in IPv6 router advertisements. Multiple prefixes can be configured in the IPv6 panel of the ASA or FWSM Add Interface or Edit Interface dialog box.

Note This dialog box is available only on ASA 7.0+ devices in routed mode; ASA 8.2+ devices in transparent mode; and FWSM 3.1+ devices in routed mode.

By default, prefixes configured as addresses on an interface are advertised in router advertisements. If you configure specific prefixes for advertisement, then only those prefixes are advertised. The valid and preferred lifetimes are counted down in real time. Alternately, a date can be set to specify the expiration of a prefix. When the expiration is reached, the prefix is no longer advertised.

Navigation Path

You can access the IPv6 Prefix Editor dialog box from the IPv6 panel of the Add Interface and Edit Interface dialog boxes: click the Add Row or Edit Row buttons beneath the table in the Interfaces IPv6 Prefixes section in either of those dialog boxes.

Related Topics

- IPv6 Address for Interface Dialog Box, page 46-52
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26

Field Reference

Table 46-14 IPv6 Prefix Editor Dialog Box

Element	Description
Address/Prefix Length	Enter an IPv6 network address, with its Prefix Length appended, where the Prefix Length integer indicates how many of the high-order, contiguous bits of the address represent the network portion of the address. A slash (/) must precede the Prefix Length. For example, 3FFE:C00:0:1::/64.
Default	If this box is checked, the settings in this dialog box will apply to all prefixes, rather than a single address. (When checked, the Address/Prefix Length field is disabled.)
No Advertisements	When checked, hosts on the local link cannot use the specified prefix in advertisements.
Off Link	When checked, the specified prefix is “off-link”; that is, not locally reachable on the link. When on-link (the default), the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link.
No Auto-Configuration	When checked, hosts on the local link cannot use the specified prefix for IPv6 autoconfiguration. When auto-configuration is on (the default), hosts on the local link can use the specified prefix for IPv6 autoconfiguration.

Table 46-14 IPv6 Prefix Editor Dialog Box (continued)

Element	Description
Prefix Lifetime	You can expand this section of the dialog box to display the following expiration options: • **Lifetime Duration** – Select this option to define prefix expiration as a length of time; the following options are enabled: – **Valid Lifetime** – The amount of time (in seconds) that the specified IPv6 prefix is advertised as being valid. Enter a value from 0 to 4294967295 seconds. The maximum value represents infinity (that is, the lifetime does not expire), which can also be specified by the checking the **Infinite** box. The default is 2592000 (30 days). – **Preferred Lifetime** – The amount of time (in seconds) that the specified IPv6 prefix is advertised as being preferred. Enter a value from 0 to 4294967295 seconds. The maximum value represents infinity (that is, the lifetime does not expire), which can also be specified by the checking the **Infinite** box. The default is 604800 (7 days). The Preferred Lifetime must be less than or equal to the Valid Lifetime. • **Lifetime Expiration Date** – Select this option to define prefix expiration as a specific date. Note that acceptable values for this date can range from today's date to one year from today's date. The following options are enabled: – **Valid** – The prefix is advertised as being valid until this date and time are reached. Enter a date in the form `Mmm dd yyyy` (that is, three-letter month abbreviation, two-digit date, and four-digit year), or click the calendar icon to select a date from a scrolling calendar. Also, enter the time of expiration on the specified date, in the form `hh:mm`, based on a 24-hour clock. – **Preferred** – The prefix is advertised as being preferred until this date and time are reached. Enter a date in the form `Mmm dd yyyy` (that is, three-letter month abbreviation, two-digit date, and four-digit year), or click the calendar icon to select a date from a scrolling calendar. Also, enter the time of expiration on the specified date, in the form `hh:mm`, based on a 24-hour clock. The Preferred date and time must be earlier than or equal to the Valid date and time.

Add or Edit DHCPv6 Pool Dialog Box

This dialog box is used to add or edit the DHCPv6 Server Pool. For clients that use StateLess Address Auto Configuration (SLAAC) in conjunction with the Prefix Delegation feature, you can configure the ASA to provide information such as the DNS server or domain name when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients.

Navigation Path

- Choose **Policy Objects** from the **Manage** menu, or click the Policy Object Manager button in the button bar, to open the Policy Object Manager pane in the lower section of the Configuration Manager window. Select **Pool Objects > DHCPv6 Pool Object** from the Object Type Selector. Right-click inside the work area and select **New Object** (and select an object type), or right-click a row and select **Edit Object**; you also can use the related buttons at the bottom of the pane to open either dialog box.

OR

- You can access the Add DHCPv6 Pool dialog box from DHCPv6 Pool Selector dialog box: click the Add Row or Edit Row buttons beneath the Available DHCPv6 Pool table. The DHCPv6 Pool Selector dialog box can been accessed from the Server Pool radio button in the Interface IPv6 DHCP section of the IPv6 panel of the Add Interface and Edit Interface dialog box.

Related Topics

- IPv6 Address for Interface Dialog Box, page 46-52
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26

Field Reference

Table 46-15 Add DHCPv6 Pool Dialog Box

Element	Description
Name	The DHCPv6 Pool name should not exceed 200 characters. Object names are not case-sensitive.

- Configure parameters on one or more tabs, to provide responses to IR messages to clients.
- For each of these tabs, specify the following as appropriate:
  - DNS/SIP/ NIS/ NISP/ SNTP Server: Enter a server name. Make sure that the IPv6 addresses are in the correct format. For more information on IPv6 address format, see http://www.ietf.org/rfc/rfc2373.txt.
  - DNS/ SIP/NIS/NISP Domain Name: Enter a domain name. Domain names must begin and end with a digit/letter, only letters, digits and hyphen are allowed as internal characters, labels are separated by a dot.Each label must be up to 63 characters and the entire host name has a maximum of 255 characters. For more information on domain names format, see http://www.ietf.org/rfc/rfc1123.txt.

Note The import command uses one or more parameters that the ASA obtained from the DHCPv6 server on the Prefix Delegation client interface. You can mix and match manually-configured parameters with imported parameters; however, you cannot configure the same parameter manually and in the import command.

Server tab	(Optional) Specify DNS Server Name and Domain Name.
SIP tab	(Optional) Specify SIP Server Name and SIP Domain Name.
NIS tab	(Optional) Specify NIS Server Name and NIS Domain Name.
NISP tab	(Optional) Specify NISP Server Name and NISP Domain Name.
SNTP tab	(Optional) Specify SNTP Server Name.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 46-15 Add DHCPv6 Pool Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Device Interface: IP Type (PIX/ASA 7.0+)

A security device operating in single-context, routed mode requires IP addressing for its interfaces; however, firewall interfaces do not have IP addresses until you assign them. Note that in transparent mode, the device acts as an access-control bridge (a "bump in the wire")—you assign different VLANs to each interface, but IP addressing is not necessary.

The Add/Edit Interface dialog box presented for an independent ASA or PIX 7.0+ device in single-context, routed mode includes the section **IP Type**, where you specify the type of IP addressing for the interface and provide related parameters, as described here. (The IP Type section of the Add/Edit Interface dialog box for PIX 6.3 devices is described in Device Interface: IP Type (PIX 6.3), page 46-30.)

In multiple-context mode, interface IP addresses are set in the context configuration.

Note Do not use addresses previously used for routers, hosts, or any other firewall device commands, such as an IP address in the global pool or a static NAT entry.

Also, do not specify IP Type information for an interface you intend to use as a redundant interface.

Step 1 In the Add/Edit Interface dialog box, choose a method for address assignment (**Static IP**, **Use DHCP**, or **PPPoE (PIX and ASA 7.2+)**) from the **IP Type** list, and then provide related parameters, as follows:

- **Static IP** – Provide a static **IP Address** and **Subnet Mask** that represents the security device on this interface's connected network. The IP address must be unique for each interface.

  The Subnet mask can be expressed in dotted decimal format (for example, 255.255.255.0), or by entering the number of bits in the network mask (for example, 24). Beginning from Version 4.13, Cisco Security Manager allows you to use 255.255.255.254 for point to point interface. Do not use 255.255.255.255 for an interface connected to the network because this will stop traffic on that interface. If you omit the Subnet Mask value, a "classful" network is assumed, as follows:

  - The Class A netmask (255.0.0.0) is assumed if the first octet of the IP Address is 1 through 126 (i.e., addresses 1.0.0.0 through 126.255.255.255).
  - The Class B netmask (255.255.0.0) is assumed if the first octet of the IP Address is 128 through 191 (i.e., addresses 128.0.0.0 through 191.255.255.255).
  - The Class C netmask (255.255.255.0) is assumed if the first octet of the IP Address is 192 through 223 (i.e., addresses 192.0.0.0 through 223.255.255.255).

  Note Do not use addresses previously used for routers, hosts, or any other firewall device commands, such as an IP address in the global pool or a static NAT entry.

- **Use DHCP** – Enables Dynamic Host Configuration Protocol (DHCP) for automatic assignment of an IP address from a DHCP server on the connected network. The following options become available:
  - **DHCP Learned Route Metric** (required) – Assign an administrative distance to the learned route. Valid values are 1 to 255. The administrative distance for learned routes defaults to 1.

    All routes have a value or "metric" that represents its priority of use. (This metric is also referred to as "administrative distance.") When two or more routes to the same destination are available, devices use administrative distance to decide which route to use.
  - **Obtain Default Route using DHCP** – Select this option to obtain a default route from the DHCP server so that you do not need to configure a default static route. See also Configuring Static Routes, page 56-128.
  - **Enable Tracking for DHCP Learned Route** – If Obtain Default Route using DHCP is selected, you can select this option to enable route tracking via a specific Service Level Agreement (SLA) monitor. The following option becomes available:
  - **Tracked SLA Monitor** – Required if Enable Tracking for DHCP Learned Route is selected. Enter or Select the name of the SLA monitor object that defines the route tracking (connectivity monitoring) to be applied to this interface. See Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8 for more information.
- **PPPoE (PIX and ASA 7.2+)** – Enables Point-to-Point Protocol over Ethernet (PPPoE) for automatic assignment of an IP address from a PPPoE server on the connected network; this option is not supported with failover. The following options become available:
  - **VPDN Group Name** (required) – Choose the Virtual Private Dialup Network (VPDN) group that contains the authentication method and user name/password to use for network connection, negotiation and authentication. See Managing VPDN Groups, page 46-72 for more information.
  - **IP Address** – If provided, this static IP address is used for connection and authentication, instead of a negotiated address.
  - **Subnet Mask** – The subnet mask to be used in conjunction with the provided IP Address.
  - **PPPoE Learned Route Metric** (required) – Assign an administrative distance to the learned route. Valid values are 1 to 255; defaults to 1.

    All routes have a value or "metric" that represents its priority of use. (This metric is also referred to as "administrative distance.") When two or more routes to the same destination are available, devices use administrative distance to decide which route to use.
  - **Obtain Default Route using PPPoE** – Select this option to obtain a default route from the PPPoE server; sets the default routes when the PPPoE client has not yet established a connection. When using this option, you cannot have a statically defined route in the configuration.
  - **Enable Tracking for PPPoE Learned Route** – If Obtain Default Route using PPPoE is selected, you can select this option to enable route tracking for PPPoE-learned routes. The following options become available:
  - **Dual ISP Interface** – If you are defining interfaces for dual ISP support, choose Primary or Secondary to indicate which connection you are configuring.
  - **Tracked SLA Monitor** – Required if Enable Tracking for DHCP Learned Route is selected. Enter or Select the name of the SLA monitor object that defines the route tracking (connectivity monitoring) to be applied to this interface. See Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8 for more information.

Note You can configure DHCP and PPPoE only on the outside interface of a firewall device. If you have already configured PPPoE on the outside interface, it is no longer available as an option.

Step 2 Continue configuring the device interface in the Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

Device Interface: MAC Address

By default, a physical interface uses its "burned-in" MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.

A redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then its MAC address changes to match the MAC address of the interface that is now listed first. If you manually assign a MAC address to the redundant interface, that is used regardless of the physical-interface MAC addresses.

Similarly, all interfaces assigned to an EtherChannel group share the same MAC address. By default, the EtherChannel uses the MAC address of the lowest-numbered member interface. However, you can manually configure a MAC address for the EtherChannel to prevent traffic disruption should the low-numbered interface be removed from the group.

You also might want to assign unique MAC addresses to subinterfaces. For example, your service provider might control access based on MAC addresses.

Further, if you use failover, you can provide a standby MAC address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Note The following options appear only on the Advanced tab of the Add Interface and Edit Interface dialog boxes presented by PIX 7.2+ and ASA 7.2+ devices.

(Optional) To manually assign a private MAC address to the current interface:

Step 1 In the Add/Edit Interface dialog box, provide the desired MAC address in the **Active MAC Address** field.

MAC addresses are provided in *H.H.H* format, where *H* is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE.

Note In some cases, you may have to press the Tab key after entering the Active MAC Address to activate the Standby MAC Address field.

Step 2 If desired, provide a **Standby MAC Address** for use with device-level failover.

If the active unit fails over and the standby unit becomes active, the new active unit begins using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 3 Continue configuring the device interface in the Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31.

Configuring Hardware Ports on an ASA 5505

The Interfaces page displayed for ASA 5505 devices presents two tabbed panels: *Hardware Ports* and *Interfaces*. The table on the Hardware Ports panel displays currently configured switch ports for the selected ASA 5505.

Use the Configure Hardware Ports dialog box to configure the switch ports on an ASA 5505, including setting the mode, assigning a switch port to a VLAN, and setting the Protected option. (The following dialog-box parameter descriptions also describe the fields in the Hardware Ports table.)

Caution The ASA 5505 does not support Spanning Tree Protocol for loop detection in the network. Therefore, you must ensure that any connection with the appliance does not end up in a network loop.

Navigation Path

You can access the Configure Hardware Ports dialog box by clicking Add Row or Edit Row on the Hardware Ports panel of the ASA 5505 Interfaces page. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.

Related Topics

- Understanding ASA 5505 Ports and Interfaces, page 46-6
- Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31

Field Reference

Table 46-16 Configure Hardware Ports Dialog Box

Element	Description
Enable Interface	Select this option to enable this switch port. You can deselect this option to disable the port, but retain its configuration information.
Isolated	Select this option to prevent this port from communicating with other isolated or "protected" switch ports on the same VLAN. You might want to prevent switch ports from communicating with each other if the devices on those ports are primarily accessed from other VLANs, if you do not need to allow intra-VLAN access, and you want to isolate the devices from each other in case of infection or other security breach. For example, if you have a DMZ that hosts three Web servers, you can isolate the Web servers from each other if you apply the Isolated option to each switch port. The inside and outside networks can both communicate with all three Web servers, and vice versa, but the Web servers cannot communicate with each other.
Hardware Port	Choose the switch port that you are configuring; all device ports are listed.

Table 46-16 Configure Hardware Ports Dialog Box (continued)

Element	Description
Mode	Choose a mode for this port: • **Access Port** – Sets the port to access mode. Each access port can be assigned to one VLAN. • **Trunk Port** – Sets the port to trunk mode using 802.1Q tagging. Trunk ports can carry multiple VLANs using 802.1Q tagging. Trunk mode is available only with the Security Plus license. Trunk ports do not support untagged packets, there is no native VLAN support, and the appliance drops all packets that do not contain a tag.
VLAN ID	Enter VLAN ID(s) for this port, according to the chosen Mode: • For Access Port mode, enter the ID of the VLAN to which this switch port is to be assigned. • For Trunk Port mode, you can enter multiple VLAN IDs, and multiple ID ranges (such as 4-8), separated by commas. **Note** For devices running operating system 7.2(2)18 or earlier, valid VLAN IDs are 1 to 1001; with version 7.2(2)19 or later, valid IDs are 1 to 4090.
Duplex	Choose a duplex option for the port: **Full**, **Half**, or **Auto**. The Auto setting is recommended, and the default. If you set Duplex to anything other than Auto for PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power.
Speed	Choose a speed for the port: **10**, **100**, or **Auto**. The Auto setting is recommended, and the default. If you set Speed to anything other than Auto for PoE ports Ethernet 0/6 or 0/7, then Cisco IP phones and Cisco wireless access points that do not support IEEE 802.3af will not be detected and supplied with power. The default Auto setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either Speed or Duplex must be set to Auto to enable Auto-MDI/MDIX for the interface. If you explicitly set both Speed and Duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled.

Add/Edit Bridge Group Dialog Box

A transparent firewall connects the same network on its inside and outside interfaces, and supports only the two interfaces per context. However, you can increase the number of interfaces available to a context through use of bridge groups. You can configure up to eight bridge groups; on an FWSM each group can contain two interfaces; on an ASA 9.6.1 each group can contain 64 interfaces.

Each bridge group connects to a separate network. Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the security appliance—traffic must exit the security appliance to be routed by an external router back to another bridge group in the security appliance.

You might want to use more than one bridge group if you do not want the overhead of security contexts, or want to maximize your use of security contexts. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration. For complete security policy separation, use security contexts with one bridge group in each context.

Starting from Cisco Security Manager 4.13, the Bridge-group Virtual Interface (BVI) feature is extended to the routed firewall mode. Routed firewalls are implemented by means of configuring bridge-groups. A user can configure up to eight bridge groups and on an ASA 9.7.1 (Cisco Security Manager 4.13) each group can contain upto 64 interfaces. On versions prior to Cisco Security Manager 4.13, a user can configure a maximum of two bridge groups; with each group containing a maximum limit of four interfaces. In addition to the BVI features supported in the transparent mode, the routed firewall mode includes support for the following additional communication modes:

- Inter BVI communication
- BVI to Data Port communication (Layer 2 to Layer 3) and vice versa

For FWSM 3.1+ and ASA 8.4.1+ devices in transparent mode, the Interfaces page displays two tabbed panels: Interfaces and Bridge Groups. The following information applies to the Bridge Groups panel and the Add/Edit Bridge Group dialog box; refer to Add/Edit Interface Dialog Box (PIX 7.0+/ASA/FWSM), page 46-31 for information about the Interfaces panel.

Navigation Path

You can access the Add/Edit Bridge Group dialog box from the Bridge Groups panel of the Interfaces page.

Related Topics

- Interfaces in Routed and Transparent Modes, page 46-5
- Bridging Support for FWSM 3.1, page 47-3
- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26

Field Reference

Table 46-17 Add/Edit Bridge Group Dialog Box

Element	Description
General Tab	
Bridge Group	Enter a name for this bridge group.

Table 46-17 Add/Edit Bridge Group Dialog Box (continued)

Element	Description
Name	Provide an identifier for this interface of up to 48 characters in length. The name should be a memorable name for the interface that relates to its use. However, if you are using failover, do not name interfaces that you are reserving for failover communications; this includes an EtherChannel intended for failover, as well as its member interfaces. Also, do not name interfaces intended for use as a member of a redundant-interface pair. Certain names are reserved for specific interfaces, in accordance with the interface naming conventions of the security appliance. As such, these reserved names enforce default, reserved security levels, as follows: • Inside – Connects to your internal network. Must be the most secure interface. • DMZ – "Demilitarized zone" attached to an intermediate interface. DMZ is also known as a perimeter network. You can name a DMZ interface any name you choose. Typically, DMZ interfaces are prefixed with "*DMZ*" to identify the interface type. • Outside – Connects to an external network or the Internet. Must be the least secure interface. Similarly, a subinterface name typically identifies its associated interface, in addition to its own unique identifier. For example, *DMZoobmgmt* could represent an out-of-band management network attached to the DMZ interface. **Note** Again, **do not name the interface if you intend to use it for failover, or as a member of a redundant interface**. See Chapter 50, "Configuring Failover" and Configuring Redundant Interfaces, page 46-8 for more information.
ID	Enter an identifier for this bridge group; can be an integer between 1 and 100.
Security Level	Assign a security level to the VLAN interface. Valid values are from 0-100; 100 is the most secure.
Available Interfaces	Choose from a list of available interfaces or VLANs to assign to this bridge group; all available interfaces are listed. **Note** Starting from ASA 9.7.1(Cisco Security Manager 4.13), a maximum of 64 interfaces are supported per bridge group.
Members in Group	Displays the number of interfaces in the current bridge group
IP Type	Select the IP type for the interface. • **Static IP** - Assign an IP address and subnet mask for the bridge-group interface. • **DHCP** - Use DHCP to obtain an IP address for the interface. – **Obtain Default Route using DHCP** - When selected, Cisco Security Manager uses the default route supplied by the DHCP server.

Table 46-17 Add/Edit Bridge Group Dialog Box (continued)

Element	Description
IP Address	Enter or Select a management IP address for the bridge group. A transparent firewall does not participate in IP routing. Thus, the only IP configuration required for a bridge group is this management IP address. This address is the source address for traffic originating on the security appliance, such as system messages or communications with AAA servers. You can also use this address for remote management access. **Note** IPv6 addresses are not supported for bridge groups.
Netmask	Network mask for the specified IP address. You can express the value in dotted decimal format (for example, 255.255.255.0) or by entering the number of bits in the network mask (for example, 24). **Note** Do not use 255.255.255.255 for an interface connected to the network because this will stop traffic on that interface.
Description	You can enter an optional description for this bridge group.
IPv6 Tab	
Enable IPv6	Check this box to enable IPv6 and configure IPv6 addresses on this bridge group. You can deselect this option to disable IPv6 on the bridge group, but retain the configuration information.
Enforce EUI-64	When selected, use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link is enforced. When this option is enabled on a bridge group, the source addresses of IPv6 packets received on the bridge group interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the interface identifier in an IPv6 packet is not in the Modified EUI-64 format, the packet is dropped and the following system log message is generated: `%PIX|ASA-3-325003: EUI-64 source address check failed.` Address format verification is performed only when a flow is created. Packets from an existing flow are not checked. Additionally, address verification can be performed only for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address. The Modified EUI-64 format interface identifier is derived from the 48-bit link-layer (MAC) address by inserting the hex number FFFE between the upper three bytes (OUI field) and the lower 3 bytes (serial number) of the link-layer address. To ensure the chosen address is from a unique Ethernet MAC address, the next-to-lowest order bit in the high-order byte is inverted (universal/local bit) to indicate the uniqueness of the 48-bit address. For example, an interface with a MAC address of 00E0.B601.3B7A would have a 64-bit interface ID of 02E0:B6FF:FE01:3B7A.

Table 46-17 Add/Edit Bridge Group Dialog Box (continued)

Element	Description
DAD Attempts	To specify the number of consecutive neighbor solicitation messages that are sent on a bridge group interface during duplicate address detection (DAD), enter a number from 0 to 600 in this field. Entering 0 disables duplicate address detection on the interface. Entering 1 configures a single transmission without follow-up transmissions; this is the default. Duplicate address detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection uses neighbor solicitation messages to verify the uniqueness of unicast IPv6 addresses. When duplicate address detection identifies a duplicate address, the state of the address is set to DUPLICATE and the address is not used. If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface and an error message similar to the following is issued: `%PIX-4-DUPLICATE: Duplicate address FE80::1 on outside` If the duplicate address is a global address of the interface, the address is not used and an error message is issued, similar to that shown previously for a duplicate link-local address. All configuration commands associated with the duplicate address remain as-configured while the state of the address is set to DUPLICATE. If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address, and all other IPv6 address associated with the interface are regenerated (that is, duplicate address detection is performed only on the new link-local address).
NS Interval	The interval between IPv6 neighbor solicitation retransmissions, in milliseconds. Valid values range from 1000 to 3600000 milliseconds; the default value is 1000 milliseconds. **Note** This value is included in all IPv6 router advertisements sent out on this interface.
Reachable Time	The amount of time, in milliseconds, within which a remote IPv6 node is considered still reachable, after initial reachability was confirmed. Valid values range from 0 to 3600000 milliseconds, the default value is 0. When 0 is used for the value, the reachable time is set as undetermined—it is up to the receiving devices to set and track reachable time. A configured time enables detection of unavailable neighbors. A shorter time allows detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.
Managed Config Flag	Whether or not to set the flag "managed-config-flag" in the IPv6 router advertisement packet.

Table 46-17 Add/Edit Bridge Group Dialog Box (continued)

Element	Description
Other Config Flag	Whether or not to set the flag "other-config-flag" in the IPv6 router advertisement packet.
Enable RA	When checked, IPv6 router advertisement transmissions are enabled on the interface. The following options are enabled: • **RA Lifetime** – The “router lifetime” value specifies how long nodes on the local link should consider the security appliance as the default router on the link. Valid values range from 0 to 9000 seconds; the default is 1800 seconds. Entering 0 indicates that the security appliance should not be considered a default router on the selected interface. Any non-zero value should not be less than the following RA Interval value. **Note** This value is included in all IPv6 router advertisements sent out on this interface. • **RA Interval** – The interval between IPv6 router advertisement transmissions on this interface. Valid values range from 3 to 1800 seconds, (or from 500 to 1800000 milliseconds if the following RA Interval in Milliseconds option is checked); the default is 200 seconds. The interval between transmissions should be less than or equal to the RA Lifetime value if it is non-zero. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value. • **RA Interval in Milliseconds** – Checking this option indicates that the provided RA Interval value is in milliseconds, rather than seconds.

Table 46-17 Add/Edit Bridge Group Dialog Box (continued)

Element	Description
Interface IPv6 Addresses	The IPv6 addresses assigned to the bridge group interface are specified in this section of the dialog box. • **Link-Local Address** – To override the link-local address that is automatically generated for the interface, enter the desired IPv6 link-local address in this field. The link-local address is composed of the link-local prefix FE80::/64 and the interface ID in Modified EUI-64 format. For example, an interface with a MAC address of 00E0.B601.3B7A would have a link-local address of FE80::2E0:B6FF:FE01:3B7A. An error will occur if another host is using the specified address. • **Enable Address Auto-Configuration** – Select this option to enable automatic configuration of IPv6 addresses on the interface using stateless autoconfiguration. The addresses are configured based on the prefixes received in Router Advertisement (RA) messages. If a link-local address has not been configured, then one is automatically generated for this interface. An error occurs if another host is already using the generated link-local address. – **Trust the DHCP Servers for default gateway**– Select this radio button to install a default route from Router Advertisements that come from a trusted source - the directly-connected network. – **Ignore trust and accept router advertisements** – Select this radio button to install a default route from Router Advertisements that come from another network. • The table in this section displays the IPv6 addresses assigned to this interface. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Address for Interface Dialog Box, page 46-52.
Interface IPv6 Prefixes	Use the table in this section to configure which IPv6 prefixes (that is, the network portion of the IPv6 addresses) are included in IPv6 router advertisements. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Prefix Editor Dialog Box, page 46-54.

Advanced Interface Settings (PIX/ASA/FWSM)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Advanced configuration options are available for interfaces on FWSMs and ASA/PIX 7.0+ devices operating in single-context mode and for ASA 9.0+ devices operating in single-context mode or multi-context mode.

These are general device-related settings; that is, they are not applied to individual interfaces.

Note The information in this section does not apply to PIX 6.3 devices, nor to security devices in multiple-context mode.

The Advanced Interface Settings dialog box includes the following elements:

- **MAC Address Auto** - Enable this option to automatically assign private MAC addresses to each shared context interface. You can also, optionally, set a user-defined prefix as part of the MAC address. The prefix is a decimal value between 0 and 65535. If you do not enter a prefix, then the ASA generates a default prefix. This prefix is converted to a 4-digit hexadecimal number. The prefix ensures that each ASA uses unique MAC addresses (using different prefix values), so you can have multiple ASAs on a network segment, for example,
- **Traffic between interfaces with same security levels** – This parameter controls communication between interfaces and subinterfaces on the same security level. If you enable same security interface communication, you can still configure interfaces at different security levels as usual. Refer to Enabling Traffic between Interfaces with the Same Security Level, page 46-70 for more information.
- **PPPoE Users button** – Click this button to open the PPPoE Users dialog box, where you can add, edit and delete PPPoE users, as described in Managing the PPPoE Users List, page 46-71. This option is available only for ASA and PIX 7.0+ devices.
- **VPDN Groups (PIX and ASA 7.2+)** – This table lists currently defined VPDN Groups. The buttons below the table are used to add, edit and delete VPDN group entries, as described in Managing VPDN Groups, page 46-72.
- **LACP System Priority (ASA 8.4.1+)** – All systems participating in EtherChannel link aggregation require a Link Aggregation Control Protocol (LACP) System Priority. The value can be 1 to 65535, with the higher number signifying lower priority. The default is 32768.

  This value is combined with the system MAC address to form the system's LACP identifier, and thus is applicable only for EtherChannel interfaces. See Configuring EtherChannels, page 46-9, for more information.

  Note Additional LACP parameters are available in the Edit Interface dialog box for individual interfaces assigned to an EtherChannel; see Editing LACP Parameters for an Interface Assigned to an EtherChannel, page 46-12, for more information.

  Note LACP System Priority is not supported in Cisco Firepower 9000 devices.

- **Static Port Priority (ASA 9.2.1+ Cluster in Spanned mode)** – Disables dynamic port priority in LACP. Some switches do not support dynamic port priority, so this parameter improves switch compatibility. Enabling static port priority enables support of 16 active spanned EtherChannel members. Without this parameter, only 8 active members and 8 standby members are supported. If you enable this parameter, then you cannot use any standby members; all members are active. This parameter is not part of the bootstrap configuration, and is replicated from the control unit to the data units.

Note If you enable Static Port Priority, 16 nodes can be part of a cluster instead of 8 nodes.

- **Director-Localization** – In Geo clustering where, multiple Data Center sites are supported, the inter-cluster round-trip time (RTT) latency is higher than intra-DC. This delay impacts the performance of applications like the VoIP media stream. Beginning with 4.13, the director localization is used to minimize the RTT latency and the delays in performance lookup messages. Enabling this option makes the flow owner and director to be in the same DC site, so that the flow owner lookup is done in local DC site, and the traffic is contended within the same site.

  Note Director-localization is not supported in Cisco Firepower 2100 Series, Firepower 4000 Series, and Firepower 9000 Series devices.

- **Enable Site Redundancy**—Beginning with 4.16, you can enable site redundancy to protect flows from a failed site. The site redundancy can be enabled only on the Control Unit and will be replicated to the data units in the cluster group. If the connection backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a failed site. Director localization and site redundancy are separate features; you can configure one or the other, or configure both.

  Note Site redundancy is not supported in Cisco Firepower 2100 Series, Firepower 4000 Series, and Firepower 9000 Series devices.

Navigation Path

You can open the Advanced Interface Settings dialog box by clicking the Advanced button at the bottom of the Interfaces page (for non-5505 ASAs, PIX 7.0+ devices, and FWSMs), or at the bottom of the Interfaces tab on the ASA 5505 Ports and Interfaces page.

Related Topics

- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26

Enabling Traffic between Interfaces with the Same Security Level

The Advanced Interface Settings (PIX/ASA/FWSM), page 46-68 dialog box presented for a single-context security device includes the “Traffic between interfaces with the same security level” drop-down list, as described in this section.

By default, interfaces or subinterfaces on the same security level cannot communicate with each other. Allowing communication between same-security interfaces provides the following benefits:

- You can configure more than 101 communicating interfaces.

  If you use different levels for each interface and do not assign any interfaces to the same security level, you can configure only one interface per level (0 to 100).

- You can allow traffic to flow freely between all same-security interfaces without access lists.

Note If you enable NAT control, you do not need to configure NAT between same-security-level interfaces.

Step 1 In the Advanced Interface Settings dialog box, choose the option that identifies how you want this device to handle **Traffic between interfaces with the same security levels**:

- **Disabled**—Communication between interfaces on the same security level is not allowed.
- **Inter-interface**—Enables traffic flows between interfaces with the same security level setting. When this option is enabled, you are not required to define translation rules to enable traffic flow between interfaces in the firewall device.
- **Intra-interface**—Enables traffic flows between subinterfaces with the same security level setting. When this option is enabled, you are not required to define translation rules to enable traffic flow between subinterfaces assigned to an interface.
- **Both**—Allows both intra- and inter-interface communications among interfaces and subinterfaces with the same security level.

Step 2 Continue with Advanced Interface Settings (PIX/ASA/FWSM), page 46-68 configuration, or click OK to close the Advanced Interface Settings dialog box.

Managing the PPPoE Users List

Point-to-Point Protocol over Ethernet (PPPoE) allows standard PPP communication between a security device and an external ISP, via an Ethernet interface on the device. To establish a communication link, the device must provide authentication credentials and obtain network parameters. This is accomplished using a Virtual Private Dialup Network (VPDN) group, which basically consists of established PPPoE user credentials (i.e., a user name and password) and an authentication protocol. See Managing VPDN Groups, page 46-72 for more information about VPDN groups.

The PPPoE user credentials available for use with VPDN groups are maintained in the PPPoE Users dialog box, which you can access from the Advanced Interface Settings (PIX/ASA/FWSM), page 46-68 dialog box, and from the Add/Edit VPND Group dialog boxes.

Adding and Editing PPPoE Users

The PPPoE Users dialog box presents a table of currently defined PPPoE users, along with standard Add Row, Edit Row, and Delete Row buttons. The Add Row button opens the Add PPPoE User dialog box; the Edit Row button opens the virtually identical Edit PPPoE User dialog box.

Enter or edit the following PPPoE user parameters, and then click OK to close the Add (Edit) PPPoE User dialog box and return to the Advanced Interface Settings dialog box.

Note PPPoE user options are not available on Firewall Service Modules (FWSMs).

Field Reference

Table 46-18 Add and Edit PPPoE User Dialog Boxes

Element	Description
Username	The name assigned to this user account; generally provided by the external ISP.
Password	The password assigned to this user account; also generally provided by the external ISP.

Table 46-18 Add and Edit PPPoE User Dialog Boxes (continued)

Element	Description
Confirm	Re-enter the password.
Store Username and Password in Local Flash	If checked, this PPPoE user information will be stored in the device's local flash memory, ensuring it cannot be inadvertently overwritten.

Managing VPDN Groups

A Virtual Private Dialup Network (VPDN) group—basically an established PPPoE user and an authentication protocol—is used by a security device to contact an external ISP and authenticate itself, in order to establish a PPPoE communications link and obtain network parameters. (See Managing the PPPoE Users List, page 46-71 for information about establishing PPPoE users.)

Available VPDN groups are maintained in the Advanced Interface Settings dialog box, which opens when you click the Advanced button at the bottom of the Interfaces page, as described in Advanced Interface Settings (PIX/ASA/FWSM), page 46-68.

Adding and Editing VPND Groups

The Advanced Interface Settings dialog box includes a table of currently defined VPDN groups, and standard Add Row, Edit Row, and Delete Row buttons. The Add Row button opens the Add VPDN Group dialog box; the Edit Row button opens the virtually identical Edit VPDN Group dialog box.

Enter or edit the following VPDN group parameters, and then click OK to close the Add (Edit) VPDN Group dialog box and return to the Advanced Interface Settings dialog box.

Note VPDN group options are not available on Firewall Service Modules (FWSMs).

Field Reference

Table 46-19 Add and Edit VPND Group Dialog Boxes

Element	Description
Group Name	A name to identify this group in Security Manager; up to 63 characters.
PPPoE Username	The name identifying the PPPoE credentials to be used by this group for authentication with an ISP; choose from the list of available PPPoE users. Choose **Edit User** from this list to open the PPPoE Users dialog box, where you can add or edit a user for this option. Refer to Managing the PPPoE Users List, page 46-71 for information about creating and editing users.
PPP Authentication	Select the PPP Authentication method: • **PAP** – Password Authentication Protocol, with exchange of credentials in clear text. • **CHAP** – Challenge Handshake Authentication Protocol, with encrypted credential exchange. • **MSCHAP** – Microsoft's CHAP, version 1 only.

VXLAN

Virtual eXtensible LANs (VXLAN) act as Layer 2 virtual networks over Layer 3 physical networks to stretch Layer 2 networks. VXLAN provides the same Ethernet Layer 2 network services as VLAN does, but with greater extensibility and flexibility. Compared to VLAN, VXLAN offers the following benefits:

- Flexible placement of multitenant segments throughout the data center.
- Higher scalability to address more Layer 2 segments—up to 16 million VXLAN segments.

Beginning with version 4.9, Security Manager supports VXLAN for ASA, ASAv, and ASASM devices on version 9.4(1) and higher.

Note VxLAN is not supported on IOS, FWSM, ASR, or ISR devices.

To configure VXLAN, follow these steps:

1. Configuring VXLAN Policy, page 46-73
2. Create a Configuring VNI Interfaces, page 46-15 and associate the configured VXLAN policy to the VNI interface.

Configuring VXLAN Policy

To configure VXLAN you must first configure VXLAN policy and then create a VNI interface and associate the configured VXLAN policy to the VNI interface. This section describes how to configure VXLAN policy.

Navigation Path

To access the VXLAN page, go to **Device View**, select an ASA, ASAv, or ASASM device and then click **VxLAN** from **Policies**.

Related Topics

- VXLAN, page 46-73
- Configuring VNI Interfaces, page 46-15

Field Reference

Table 46-20 VxLAN

Element	Description
Enable VXLAN Port Number VXLAN Destination Port	Check this box if you want to change the value of the VXLAN Destination Port from the default 4789. If checked, enter a numeric value in the range from 1024 to 65535.
Network Virtualization Endpoint (NVE)	
Enable NVE	When selected, it enables you to select the VTEP Tunnel Interface.
VXLAN NVE No	The value of VXLAN NVE Number is "1". You cannot edit this value.
Enable NVE Encapsulation	Select this checkbox to enable NVE Encapsulation using VXLAN.
VTEP Tunnel Interface	Click Select and choose the VTEP Tunnel Interface.

Table 46-20 VxLAN (continued)

Element	Description
Enable VTEP IP Address Or Multicast Traffic Address	Select one of the following: • Peer VTEP IP Address—Manually specify the peer VTEP IP address. If you specify the peer IP address, you cannot use multicast group discovery. Multicast is not supported in multiple context mode. You can only specify one peer for the VTEP. Note that the peer VTEP IP address must be reachable from the VTEP Tunnel Interface, else deployment will fail. If you have used peer IP address in VXLAN policy, you cannot configure multicast IP address on the Interface page, including the VNI interface. • Default Multicast IP Address—Specify a default multicast group for all associated VNI interfaces. The range of IP addresses is from 224.0.0.0 to 239.255.255.255. If you do not configure the multicast group per VNI interface, then this group is used. If you configure a group at the VNI interface level, then that group overrides this setting.
Save	Click Save to save the VXLAN settings.

Configuring Bridging Policies on Firewall Devices

This chapter contains the following topics:

- About Bridging on Firewall Devices, page 47-1
- Bridging Support for FWSM 3.1, page 47-3
- ARP Table Page, page 47-3
- ARP Inspection Page, page 47-5
- Managing the IPv6 Neighbor Cache, page 47-7
- MAC Address Table Page, page 47-8
- MAC Learning Page, page 47-9
- Management IP Page, page 47-10
- Management IPv6 Page (ASA 5505), page 47-11

About Bridging on Firewall Devices

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 device that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. The security appliance connects the same network on its inside and outside ports, acting as an access-control bridge; you assign different VLANs to each interface, and IP addressing is not used.

Thus, you can easily introduce a transparent firewall into an existing network—IP re-addressing is unnecessary—and maintenance is facilitated because there are no complicated routing patterns to troubleshoot and no NAT configuration.

Although the transparent-mode device acts as a bridge, Layer 3 traffic, such as IP traffic, cannot pass through the security appliance unless you explicitly permit it with specific access rules. The only traffic allowed through a firewall without an access list is ARP traffic, which you can control using ARP inspection, and IPv6 neighbor discovery.

When the security appliance runs in transparent mode, the outgoing interface of a packet is determined by performing a MAC address lookup instead of a route lookup. Route statements can still be configured, but they apply only to security appliance-originated traffic. For example, if your syslog server is located on a remote network, you must use a static route so the security appliance can reach that subnet.

Starting from Cisco Security Manager 4.13, the Bridge-group Virtual Interface (BVI) feature is extended to the routed firewall mode. Routed firewalls are implemented by means of configuring bridge-groups. A user can configure up to eight bridge groups and on an ASA 9.7.1 (Cisco Security Manager 4.13) each group can contain upto 64 interfaces. On versions prior to Cisco Security Manager 4.13, a user can configure a maximum of two bridge groups; with each group containing a maximum limit of four interfaces. In addition to the BVI features supported in the transparent mode, the routed firewall mode includes support for the following additional communication modes:

- Inter BVI communication
- BVI to Data Port communication (Layer 2 to Layer 3) and vice versa

To configure a transparent firewall, use the following policies. When configuring an ASA/PIX/FWSM device in multiple-context mode, configure these policies on each transparent security context.

- **Firewall > Access Rules**—Access rules control layer 3 and higher traffic using extended access control lists. In routed mode, some types of traffic cannot pass through the security appliance even if you allow it in an access list. For example, you can establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on access rules. Likewise, protocols like HSRP or VRRP can pass through the security appliance. However, the transparent-mode security appliance does not pass CDP packets.

  For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can provide those functions. For example, by using access rules, you can allow DHCP traffic to pass (instead of the unsupported DHCP relay feature), or multicast traffic such as that created by IP/TV.

  For more information, see Understanding Access Rules, page 16-1 and Configuring Access Rules, page 16-7.

- **Firewall > Transparent Rules**—Transparent rules control non-IP layer 2 traffic using Ethertype access control lists. For example, you can configure rules to allow AppleTalk, IPX, BPDUs, and MPLS to pass through the device. For more information, see Configuring Transparent Firewall Rules, page 23-1.

- **Platform > Bridging > ARP Table**, **ARP Inspection** and **IPv6 Neighbor Cache**—Use these policies to control the types of ARP and IPv6 traffic allowed through the bridge. If desired, you can configure static ARP and IPv6 neighbor cache entries and drop any traffic not defined by those static rules. Enable ARP inspection so that if a mismatch between the MAC address, the IP address, or the interface occurs, the security appliance drops the packet. This helps prevent ARP spoofing. For more information, see ARP Table Page, page 47-3 and ARP Inspection Page, page 47-5.

  Note The ARP Table and IPv6 Neighbor Cache are the only bridging policies available for non-transparent ASA/PIX/FWSM devices.

- **Platform > Bridging > MAC Address Table** and **MAC Learning**—Use these policies to configure static MAC-IP address mappings and to enable or disable MAC learning. MAC learning is enabled by default, which allows the appliance to add MAC-IP address mappings as traffic passes through the interface. If you want to prevent all traffic except from static entries, you can disable MAC learning. For more information, see MAC Address Table Page, page 47-8 and MAC Learning Page, page 47-9.

- **Platform > Bridging > Management IP** and **Platform > Bridging > Management IPv6**—Use these policies to configure a management IP address that Security Manager can use to communicate with the device.

Note The Management IP and Management IPv6 pages are not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

If you change the management IP address, you also need to update the device properties for the device or security context. Follow these steps:

- Change the management IP address, save and submit your changes.
- Deploy your changes to the device.
- In Device view, select the device or security context, then select **Tools > Device Properties**. On the General page, enter the new management IP address in the IP Address field. On the Credentials tab, update the username and password fields with account credentials that can log into the management interface. Security Manager will now use this address and user account for subsequent deployments and device communication.

For more information, see Management IP Page, page 47-10.

Related Topics

- Bridging Support for FWSM 3.1, page 47-3
- Interfaces in Routed and Transparent Modes, page 46-5
- Transparent Rules Page, page 23-3

Bridging Support for FWSM 3.1

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

Although FWSM 3.1 can support multiple L2 interface pairs, Security Manager lets you specify no more than two L2 interfaces (a single interface pair), and one associated management IP address. That means only one bridge group with two named interfaces associated is provisioned with a management IP address. If the device configuration contains a maximum of one bridge group and two named interfaces, it is valid for discovery. All other scenarios result in an error message and the commands are ignored during discovery. Furthermore, discovery will not show any bridge-group information in Security Manager, although the bridge-group commands will be generated during deployment. Bridge group 1 will be deployed and used in transparent rule policies if no bridge group exists in the device configuration.

Related Topics

- About Bridging on Firewall Devices, page 47-1

ARP Table Page

Use the ARP Table page to add static ARP entries that map a MAC address to an IP address and identifies the interface through which the host is reached.

Navigation Path

- (Device view) Select **Platform > Bridging > ARP Table** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > ARP Table** from the Policy Type selector. Right-click **ARP Table** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Add/Edit ARP Configuration Dialog Box, page 47-5
- About Bridging on Firewall Devices, page 47-1
- ARP Inspection Page, page 47-5
- MAC Address Table Page, page 47-8
- MAC Learning Page, page 47-9
- Management IP Page, page 47-10

Field Reference

Table 47-1 ARP Table Page

Element	Description
Timeout (seconds)	The amount of time, between 60 and 4294967 seconds, before the security appliance rebuilds the ARP table. The default is 14400 seconds. Rebuilding the ARP table automatically updates new host information and removes old host information. You might want to reduce the timeout because the host information changes frequently. **Note** The timeout applies to the *dynamic* ARP table, and not the static entries contained in the ARP table.
ARP Table	
Interface	The interface to which the host is attached.
IP Address	The IP address of the host.
MAC Address	The MAC address of the host.
Alias Enabled	Indicates whether the security appliance performs proxy ARP for this mapping. If this setting is enabled and the security appliance receives an ARP request for the specified IP address, it responds with the security appliance MAC address. When the security appliance receives traffic destined for the host belonging to the IP address, the security appliance forwards the traffic to the host MAC address that you specify in this command. This feature is useful if you have devices that do not perform ARP, for example. **Note** In transparent firewall mode, this setting is ignored and the security appliance does not perform proxy ARP.

Add/Edit ARP Configuration Dialog Box

Use the Add/Edit ARP Configuration dialog box to add a static ARP entry that maps a MAC address to an IP address and identifies the interface through which the host is reached.

Navigation Path

You can access the Add/Edit ARP Configuration dialog box from the ARP Table page. For more information about the ARP Table page, see ARP Table Page, page 47-3.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3

Field Reference

Table 47-2 Add/Edit ARP Configuration dialog box

Element	Description
Interface	The name of the interface to which the host network is attached.
IP Address	The IP address of the host.
MAC Address	The MAC address of the host; for example, 00e0.1e4e.3d8b.
Enable Alias	When selected, enables proxy ARP for this mapping. If the security appliance receives an ARP request for the specified IP address, it responds with the security appliance MAC address. When the security appliance receives traffic destined for the host belonging to the IP address, the security appliance forwards the traffic to the host MAC address that you specify in this command. This feature is useful if you have devices that do not perform ARP, for example. **Note** In transparent firewall mode, this setting is ignored and the security appliance does not perform proxy ARP.

ARP Inspection Page

Use the ARP Inspection page to configure ARP inspection for a transparent firewall. ARP inspection is used to prevent ARP spoofing.

Navigation Path

- (Device view) Select **Platform > Bridging > ARP Inspection** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > ARP Inspection** from the Policy Type selector. Right-click **ARP Inspection** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Add/Edit ARP Inspection Dialog Box, page 47-6
- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3

- MAC Address Table Page, page 47-8
- MAC Learning Page, page 47-9
- Management IP Page, page 47-10

Field Reference

Table 47-3 ARP Inspection Page

Element	Description
ARP Inspection Table	
Interface	The name of the interface to which the ARP inspection setting applies.
ARP Inspection Enabled	Indicates whether ARP inspection is enabled on the specified interface.
Flood Enabled	Indicates whether packets that do not match any element of a static ARP entry should be flooded out all interfaces except the originating interface. If there is a mismatch between the MAC address, the IP address, or the interface, the security appliance drops the packet. If you do not select this check box, all non-matching packets are dropped. **Note** The dedicated management interface, if present, never floods packets even if this parameter is set to flood.

Add/Edit ARP Inspection Dialog Box

Use the Add/Edit ARP Inspection dialog box to enable or disable ARP inspection for a transparent firewall interface.

Navigation Path

You can access the Add/Edit ARP Inspection dialog box from the ARP Inspection page. For more information about the ARP Inspection page, see ARP Inspection Page, page 47-5.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- ARP Inspection Page, page 47-5

Field Reference

Table 47-4 Add/Edit ARP Inspection dialog box

Element	Description
Interface	The name of the interface for which you are enabling or disabling ARP inspection.
Enable ARP Inspection on this interface	When selected, enables ARP inspection on the specified interface.

Table 47-4 Add/Edit ARP Inspection dialog box (continued)

Element	Description
Flood ARP packets	When selected, packets that do not match any element of a static ARP entry are flooded out all interfaces except the originating interface. If there is a mismatch between the MAC address, the IP address, or the interface, the security appliance drops the packet. If you do not select this check box, all non-matching packets are dropped. **Note** The dedicated management interface, if present, never floods packets even if this parameter is set to flood.

Managing the IPv6 Neighbor Cache

Use the IPv6 Neighbor Cache page to manage static IPv6 neighbor entries that map a MAC address to an IPv6 address, and identify the interface through which the neighbor host is reached, to provide address-resolution functions for IPv6. This is available on ASA 7.0+ devices only.

Note The IPv6 Neighbor Cache entries are the IPv6 equivalent of the static ARP entries, managed on the ARP Table Page, page 47-3.

If an entry for a specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry. Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.

The IPv6 Neighbor Cache page is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons. (These are standard buttons, as described in Using Tables, page 1-48.) The Add Row button opens the Add IPv6 Neighbor Cache Configuration dialog box, and Edit Row opens the Edit IPv6 Neighbor Cache Configuration dialog box. Other than the titles, the two dialog boxes are identical.

Note Ensure that IPv6 is enabled on at least one interface before trying to add a neighbor.

Field Reference

Table 47-5 Add/Edit IPv6 Neighbor Cache Configuration dialog boxes

Element	Description
Interface	Enter or Select the name of the interface on which to add the neighbor.
IP Address	Enter the IPv6 address that corresponds to the local data-link address. (If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry.)
MAC Address	Enter the local data-line (hardware) MAC address of the host; for example, 00e0.1e4e.3d8b.

MAC Address Table Page

Use the MAC Address Table page to add static MAC address entries to the MAC Address table. The table associates the MAC address with the source interface so that the security appliance knows to send any packets addressed to the device out the correct interface.

Navigation Path

- (Device view) Select **Platform > Bridging > MAC Address Table** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > MAC Address Table** from the Policy Type selector. Right-click **MAC Address Table** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Add/Edit MAC Table Entry Dialog Box, page 47-8
- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3
- ARP Inspection Page, page 47-5
- MAC Learning Page, page 47-9
- Management IP Page, page 47-10

Field Reference

Table 47-6 MAC Address Table Page

Element	Description
Aging Time (minutes)	Sets the number of minutes, between 5 and 720 (12 hours), that a MAC address entry stays in the MAC address table before timing out. 5 minutes is the default.
MAC Address Table	
Interface	The interface to which the MAC address is associated.
MAC Address	The MAC address; for example, 00e0.1e4e.3d8b.

Add/Edit MAC Table Entry Dialog Box

Use the Add/Edit MAC Table Entry dialog box to add static MAC address entries to the MAC Address table or to modify entries in the MAC Address table.

Navigation Path

You can access the Add/Edit MAC Table Entry dialog box from the MAC Address Table page. For more information about the MAC Address Table page, see MAC Address Table Page, page 47-8.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- MAC Address Table Page, page 47-8

Field Reference

Table 47-7 Add/Edit MAC Table Entry dialog box

Element	Description
Interface	The interface to which the MAC address is associated.
MAC Address	The MAC address; for example, 00e0.1e4e.3d8b.

MAC Learning Page

Use the MAC Learning page to enable or disable MAC address learning on an interface. By default, each interface learns the MAC addresses of entering traffic, and the security appliance adds corresponding entries to the MAC address table. You can disable MAC address learning if desired; however, unless you statically add MAC addresses to the table, no traffic can pass through the security appliance.

Navigation Path

- (Device view) Select **Platform > Bridging > MAC Learning** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > MAC Learning** from the Policy Type selector. Right-click **MAC Learning** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Add/Edit MAC Learning Dialog Box, page 47-9
- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3
- ARP Inspection Page, page 47-5
- MAC Address Table Page, page 47-8
- Management IP Page, page 47-10

Field Reference

Table 47-8 MAC Learning Page

Element	Description
MAC Learning Table	
Interface	The interface to which the MAC learning setting applies.
MAC Learning Enabled	Indicates whether the security appliance learns MAC addresses from traffic entering the interface.

Add/Edit MAC Learning Dialog Box

Use the Add/Edit MAC Learning dialog box to enable or disable MAC address learning on an interface.

Navigation Path

You can access the Add/Edit MAC Learning dialog box from the MAC Learning page. For more information about the MAC Learning page, see MAC Learning Page, page 47-9.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- MAC Learning Page, page 47-9

Field Reference

Table 47-9 Add/Edit MAC Learning dialog box

Element	Description
Interface	The interface to which the MAC learning setting applies.
MAC Learning Enabled	When selected, the security appliance learns MAC addresses from traffic entering the interface.

Management IP Page

A transparent firewall does not participate in IP routing. The only IP configuration required for the device is specification of a management IP address, which is used as the source address for traffic originating on the device, such as system messages or communications with AAA servers. You can also use this address for remote-management access.

For IPv4 traffic, the management IP address is required to pass any traffic.

Note In addition to the management IP address for the device, you can configure an IP address for the Management 0/0 or 0/1 management-only interface. This IP address can be on a separate subnet from the main management IP address.

Use the Management IP page to set the management IP address for a security device, or for a context in transparent firewall mode.

Navigation Path

- (Device view) Select **Platform > Bridging > Management IP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > Management IP** from the Policy Type selector. Right-click **Management IP** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3
- ARP Inspection Page, page 47-5
- MAC Address Table Page, page 47-8
- MAC Learning Page, page 47-9

Field Reference

Table 47-10 Management IP Page

Element	Description
Management IP Address	The management IP address.
Subnet Mask	The subnet mask that corresponds to the management IP address.

Management IPv6 Page (ASA 5505)

A transparent firewall does not participate in IP routing. The only IP configuration required for the device is specification of a management IP address, which is used as the source address for traffic originating on the device, such as system messages or communications with AAA servers. You can also use this address for remote-management access.

For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations. If you configure the global address, a link-local address is automatically configured on each interface, so you do not also need to specifically configure a link-local address. However, if you do not configure a global management address, you need to configure interface link-local addresses, as described in Configuring IPv6 Interfaces (ASA/FWSM), page 46-47. Note that you can configure both IPv6 and IPv4 management addresses on a device.

On an ASA 5505 in transparent mode, use the Management IPv6 page to enable IPv6, configure neighbor solicitation, and manage IPv6 interface addresses.

Note This page is available only on ASA 5505 version 8.2 and 8.3 devices in transparent mode.

Navigation Path

- (Device view) Select **Platform > Bridging > Management IPv6** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Bridging > Management IPv6** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- About Bridging on Firewall Devices, page 47-1
- ARP Table Page, page 47-3
- ARP Inspection Page, page 47-5
- MAC Address Table Page, page 47-8
- MAC Learning Page, page 47-9

Field Reference

Table 47-11 Management IPv6 Page

Element	Description
Enable IPv6	Check this box to enable IPv6 and configure IPv6 management-interface addresses. You can deselect this option to disable IPv6, but retain the configuration information.

Table 47-11 Management IPv6 Page (continued)

Element	Description
DAD Attempts	To specify the number of consecutive neighbor solicitation messages that are sent on an interface during duplicate address detection (DAD), enter a number from 0 to 600 in this field. Entering 0 disables duplicate address detection. Entering 1 configures a single transmission without follow-up transmissions; this is the default. Duplicate address detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection uses neighbor solicitation messages to verify the uniqueness of unicast IPv6 addresses. When duplicate address detection identifies a duplicate address, the state of the address is set to DUPLICATE and the address is not used. If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface and an error message similar to the following is issued: `%PIX-4-DUPLICATE: Duplicate address FE80::1 on outside` If the duplicate address is a global address of the interface, the address is not used and an error message is issued, similar to that shown previously for a duplicate link-local address. All configuration commands associated with the duplicate address remain as-configured while the state of the address is set to DUPLICATE. If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address, and all other IPv6 address associated with the interface are regenerated (that is, duplicate address detection is performed only on the new link-local address).
NS Interval	The interval between IPv6 neighbor solicitation retransmissions, in milliseconds. Valid values range from 1000 to 3600000 milliseconds; the default value is 1000 milliseconds.
Reachable Time	The amount of time, in milliseconds, within which a remote IPv6 node is considered still reachable, after initial reachability was confirmed. Valid values range from 0 to 3600000 milliseconds, the default value is 0. When 0 is used for the value, the reachable time is set as undetermined—it is up to the receiving devices to set and track reachable time. A configured time enables detection of unavailable neighbors. A shorter time allows detecting unavailable neighbors more quickly; however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.
Interface IPv6 Addresses	The IPv6 address(es) assigned to the management interface are listed in this table. Use the Add Row, Edit Row, and Delete Row buttons below this table to manage these entries. (These are standard buttons, as described in Using Tables, page 1-48.) Add Row and Edit Row open the IPv6 Address for Interface Dialog Box, page 46-52.

Configuring Device Administration Policies on Firewall Devices

The Device Admin section contains pages for configuring device administration policies for firewall devices.

This chapter contains the following topics:

- About AAA on Security Devices, page 48-1
- Configuring Banners, page 48-9
- Configuring Boot Image/Configuration Settings, page 48-10
- Configuring CLI Prompt, page 48-12
- Setting the Device Clock, page 48-14
- Enabling/Disabling FIPS, page 48-15
- Configuring Umbrella Global Policy, page 48-16
- Configuring Device Credentials, page 48-17
- Managing Mount Points, page 48-18
- IP Client, page 48-20
- App Agent, page 48-21

About AAA on Security Devices

Authentication-Authorization-Accounting (AAA) enables the security appliance to determine who a user is (authentication), what the user can do (authorization), and what the user did (accounting). You can use authentication alone, or with authorization and accounting. Authorization always requires a user to be authenticated first. You also can use accounting alone, or with authentication and authorization.

Authentication-Authorization-Accounting provides an extra level of protection and control for user access beyond access lists alone. For example, you can create an ACL that allows all outside users to access Telnet on a server on the DMZ network. If you want to limit user access to the server when you may not always know the IP addresses of these users, you can enable AAA to allow only authenticated and/or authorized users to make it through the security appliance. (The Telnet server enforces authentication, too; the security appliance prevents unauthorized users from attempting to access the server.)

- **Authentication**—Authentication grants access based on user identity. Authentication establishes user identity by requiring valid user credentials, which are typically a user name and password. You can configure the security appliance to authenticate the following items:
  - Administrative connections to the security appliance using Telnet, SSH, HTTPS/ASDM, or serial console.
  - The **enable** command.
- **Authorization**—Authorization controls user capabilities after users are authenticated. Authorization controls the services and commands available to each authenticated user. If you do not enable authorization, authentication alone would provide the same access to services for all authenticated users.

  If you need the control that authorization provides, you can configure a broad authentication rule, and then have a detailed authorization configuration. For example, you might authenticate inside users who attempt to access any server on the outside network, and then use authorization to limit the outside servers that a particular user can access.

  The security appliance caches the first 16 authorization requests per user, so if the user accesses the same services during the current authentication session, the security appliance does not resend the request to the authorization server.
- **Accounting**—Accounting tracks traffic that passes through the security appliance, providing a record of user activity. If you enable authentication for that traffic, you can account for traffic per user. If you do not authenticate the traffic, you can account for traffic per IP address. Accounting information includes when sessions start and stop, user name, the number of bytes that pass through the security appliance for the session, the service used, and the duration of each session.

Preparing for AAA

AAA services depend upon the use of the Local database or at least one AAA server. You can also use the Local database as a fallback for most services provided by an AAA server. Before you implement AAA, you should configure the Local database and configure AAA server groups and servers.

Configuration of the Local database and AAA servers depends upon the AAA services you want the security appliance to support. Regardless of whether you use AAA servers, you should configure the Local database with user accounts that support administrative access, to prevent accidental lock-outs and, if so desired, to provide a fallback method when AAA servers are unreachable. For more information, see Configuring User Accounts, page 51-7.

The following table provides a summary of AAA service support by each AAA server type and by the Local database. You manage the Local database by configuring user accounts on the **Platform > Device Admin > User Accounts** page (see Configuring User Accounts, page 51-7). You establish AAA server groups and add individual AAA servers to the server groups using the **Platform > Device Admin > AAA** page.

Table 48-1 Summary of AAA Support

AAA Service	Database Type							
	Local	**RADIUS**	**TACACS+**	**SDI**	**NT**	**Kerberos**	**LDAP**	**HTTP Form**
Authentication of...								
VPN users	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes[1]

Table 48-1 Summary of AAA Support (continued)

AAA Service	Database Type							
	Local	RADIUS	TACACS+	SDI	NT	Kerberos	LDAP	HTTP Form
Firewall sessions	Yes	Yes	Yes	No	No	No	No	No
Administrators	Yes	Yes	Yes	No	No	No	No	No
Authorization of...								
VPN users	Yes	Yes	No	No	No	No	Yes	No
Firewall sessions	No	Yes[2]	Yes	No	No	No	No	No
Administrators	Yes[3]	No	Yes	No	No	No	No	No
Accounting of...								
VPN connections	No	Yes	Yes	No	No	No	No	No
Firewall sessions	No	Yes	Yes	No	No	No	No	No
Administrators	No	Yes	Yes	No	No	No	No	No

[1] HTTP Form protocol supports single sign-on authentication for WebVPN users only.

[2] For firewall sessions, RADIUS authorization is supported with user-specific ACLs only, which are received or specified in a RADIUS authentication response.

[3] Local command authorization is supported by privilege level only.

Local Database

The security appliance maintains a Local database that you can populate with user accounts, which contain, at a minimum, a user name. Typically, you assign a password and a privilege level to each user name, although passwords are optional. You can manage Local user accounts on the **Platform > Device Admin > User Accounts** page (see Configuring User Accounts, page 51-7).

If you enable command authorization using the Local database, the security appliance refers to the assigned user privilege level to determine what commands are available. By default, all commands are assigned either privilege level 0 or level 15.

Note If you add users to the Local database with access to the CLI and whom you do not want to enter privileged mode, you should enable command authorization. Without command authorization, users can access privileged mode (and all commands) at the CLI using their own password if their privilege level

is 2 or greater (2 is the default). Alternatively, you can use RADIUS or TACACS+ authentication for console access so the user will not be able to use the login command, or you can set all local users to level 1 so you can control who can use the system enable password to access privileged mode.

You cannot use the local database for network access authorization.

The user accounts in the Local database can provide fallback support for console and enable-password authentication, for command authorization, and for VPN authentication and authorization. This behavior is designed to help you prevent accidental lock-out from the security appliance.

For users who need fallback support, we recommend that their user names and passwords in the Local database match their user names and passwords on the AAA servers. This provides transparent fallback support. Because the user cannot determine whether a AAA server or the local database is providing the service, using user names and passwords on AAA servers that are different than the user names and passwords in the Local database means that the user cannot be certain which user name and password should be given.

For multiple-context mode, you can configure user names in the system execution space to provide individual logins at the CLI using the **login** command; however, you cannot configure any **aaa** commands that use the local database in the system execution space.

Note VPN functions are not supported in multiple mode.

AAA for Device Administration

You can authenticate all administrative connections to the security appliance, including:

- Telnet
- SSH
- Serial console
- ASDM
- VPN management access

You can also authenticate administrators who attempt to enter enable mode. You can authorize administrative commands. You can have accounting data for administrative sessions and for commands issued during a session sent to an accounting server.

You can configure AAA for device administration using the **Platform > Device Admin > AAA** page (see About AAA on Security Devices, page 48-1).

AAA for Network Access

You can configure rules for authenticating, authorizing, and accounting for traffic passing through the firewall by using the **Firewall > AAA Rules** page (see Chapter 15, “Managing Firewall AAA Rules”). The rules you create are similar to access rules, except that they specify whether to authenticate, authorize, or perform accounting for the traffic defined; and which AAA server group the security appliance is to use to process the AAA service request.

AAA for VPN Access

AAA services for VPN access include the following:

- User account settings for assigning users to VPN groups, configured on the **Platform > Device Admin > User Accounts** page (see Configuring User Accounts, page 51-7).
- VPN group policies that can be referenced by many user accounts or tunnel groups, configured on the **Remote Access VPN > RA VPN Policies > User Group Policy** or **Site to Site VPN > User Group Policy** page.
- Tunnel group policies, configured on the **Remote Access VPN > RA VPN Policies > PIX7.0/ASA Tunnel Group Policy** or **Site to Site VPN > PIX7.0/ASA Tunnel Group Policy** page.

Configuring AAA - Authentication Tab

The AAA page presents three tabbed panels; the **Authentication** panel is presented when you navigate to the AAA page. Use these options to control privileged access to the device console, to restrict access by connection type, and to define access messages.

Use the Authorization Tab, page 48-7 to control the services and commands available to authenticated users.

Use the Accounting Tab, page 48-8 to activate tracking of console traffic, providing a record of user activity.

Navigation Path

- (Device view) Select **Platform > Device Admin > AAA** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > AAA** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- About AAA on Security Devices, page 48-1
- Configuring User Accounts, page 51-7

Using the Authentication Tab

Use the Authentication tab to enable authentication for administrator access to the security appliance. The Authentication tab also lets you configure the prompts and messages a user sees when authenticated by an AAA server.

The device will prompt for a user name and password before you can enter commands. If the authentication server is offline, wait until the console login request times out. You can then access the console with the firewall username and the enable password.

Field Reference

Table 48-2 Authentication Tab

Element	Description
Require AAA Authentication to allow use of privileged commands	
Enable	Requires authentication from an AAA server to allow a user to access EXEC mode on the firewall. This option allows up to three attempts to access the firewall console. If this number is exceeded, an “access denied” message is displayed. When checked, the Server Group field is enabled.

Table 48-2 Authentication Tab (continued)

Element	Description
Server Group	Enter or Select the name of an AAA server to contact for user authentication.
Use LOCAL when server group fails	Check this box to use the LOCAL database as back-up if the selected server fails. (This option is not enabled until you provide a Server Group.)
Require AAA Authentication for the following types of connections	
Select the connections that require authentication. For each type, users are allowed up to three attempts to access the firewall console. If this number is exceeded, an "access denied" message is displayed. Select each connection option individually: • **HTTP** – Require AAA authentication when a user initiates an HTTPS connection to the firewall console. • **Serial** – Require AAA authentication when a user initiates a connection to the firewall console via the serial console cable. • **SSH** – Require AAA authentication when a user initiates a Secure Shell (SSH) connection to the console. • **Telnet** – Require AAA authentication when a user initiates a Telnet connection to the firewall console. For each selected connection, provide a Server Group and indicate whether the LOCAL database is used as a back-up: • **Server Group** – Enter or Select the name of an AAA server to contact for user authentication. • **Use LOCAL when server group fails** – Check this box to use the LOCAL database as back-up if the selected server fails. (This option is not enabled until you provide a Server Group.)	
Authentication Prompts	
Login Prompt	Enter the prompt a user will see when logging in to the security appliance.
Accepted Message	Enter the message displayed when successfully authenticated.
Rejected Message	Enter the message displayed when authentication fails for any reason.
Rejected Message for Invalid Credentials	Enter the message displayed when authentication fails following entry of unknown or invalid credentials. Available only on FWSM 3.2+ devices.
Rejected Message for Expired Password	Enter the message displayed when authentication fails following entry of an expired password. Available only on FWSM 3.2+ devices.
Maximum Local Authentication Failed Attempts	Specify the number of times the device will try to authenticate a user in the LOCAL database before that account is locked; valid values are 1 through 16. Available only on ASA/PIX 7.01+ and FWSM 3.11+ devices.

Table 48-2 Authentication Tab (continued)

Element	Description
Login History	Check this to enable the login history reporting feature. When enabled, information about all the administrative login attempts is collected and displayed on the ASA, immediately after a successful login. This includes the following information: • Date and time of the last login • Location of last login (terminal or IP address) • Number of unsuccessful login attempts since the last successful login. • Number of successful login attempts occurring during an organization-defined time period. **Note** This feature is enabled by default.
Duration (Optional)	Enter the number of days for which login events will be saved. When no value is specified here, the login history is unbounded. **Note** The default value is 90 days.

Authorization Tab

The Authorization tab allows you to configure authorization for accessing firewall commands.

Navigation Path

You can access the Authorization tab from the AAA page; see Configuring AAA - Authentication Tab, page 48-5.

Related Topics

- About AAA on Security Devices, page 48-1
- Accounting Tab, page 48-8

Field Reference

Table 48-3 Authorization Tab

Element	Description
Enable Authorization for Command Access	Requires authorization for accessing firewall commands.
Server Group	Specify the server group to use for authorization.
Use LOCAL when server group fails	Uses the LOCAL server group if the selected server group fails.

Table 48-3 Authorization Tab (continued)

Element	Description
Enable Authorization for exec shell access (ASA 8.0(2)+ only)	When selected, enables management authorization. After enabling management authorization, specify whether to use the remote server or the local database for authorization: • **Local Server**—the local user database is the source for the username entered and the Service-Type and Privilege-Level attributes assigned. • **Remote Server**—the same server is used for both authentication and authorization.
Auto Enable Authorization for exec shell access (ASA 9.1(5)+ only)	Allows users with sufficient privileges from the login authentication server to be placed directly in privileged EXEC mode. Otherwise, users are placed in user EXEC mode. These privileges are determined by the Service-Type and Privilege-Level attributes that are required to enter each EXEC mode. To enter privileged EXEC mode, users must have a Service-Type attribute of Administrative and a Privilege Level attribute of greater than 1 assigned to them. This option is not supported in the system context. However, if you configure Telnet or serial authentication in the admin context, then authentication also applies to sessions from the switch to the ASASM.
Enable Authorization for HTTP Connection	When selected, authorization via HTTP is enabled. Authorization of username is disabled by default.
Server Group	Select the server group to use for authorization.
Use LOCAL when server group fails (ASA 9.4(1)+ only)	Uses the LOCAL server group if the selected server group fails.

Accounting Tab

Use the Accounting tab to enable accounting for access to the firewall device and for access to commands on the device.

Navigation Path

You can access the Accounting tab from the AAA page; see Configuring AAA - Authentication Tab, page 48-5.

Related Topics

- About AAA on Security Devices, page 48-1
- Authorization Tab, page 48-7

Field Reference

Table 48-4 Accounting Tab

Element	Description
Require AAA Accounting for privileged commands	
Enable	When selected, enables the generation of accounting records to mark the entry to and exit from privileged mode for administrative access via the console.
Server Group	Specify the server or group of RADIUS or TACACS+ servers to which accounting records are sent.
Require AAA Accounting for the following types of connections	
Connection type	Specify the connection types that will generate accounting records: • **HTTP**—Enable or disable the generation of accounting records to mark the establishment and termination of admin sessions created over HTTP. Valid server group protocols are RADIUS and TACACS+. • **Serial**—Enable or disable the generation of accounting records to mark the establishment and termination of admin sessions that are established via the serial interface to the console. Valid server group protocols are RADIUS and TACACS+. • **SSH**—Enable or disable the generation of accounting records to mark the establishment and termination of admin sessions created over SSH. Valid server group protocols are RADIUS and TACACS+. • **Telnet**—Enable or disable the generation of accounting records to mark the establishment and termination of admin sessions created over Telnet. Valid server group protocols are RADIUS and TACACS+.
Server Group	Specify the server or group of RADIUS or TACACS+ servers to which accounting records are sent.
Require Accounting for command access	
Enable	When selected, enables the generation of accounting records for commands entered by an administrator/user.
Server Group	Provides a drop-down menu from which you can choose the server or group of RADIUS or TACACS+ servers to which accounting records are sent.
Privilege Level	Minimum privilege level that must be associated with a command for an accounting record to be generated. The default privilege level is 0.

Configuring Banners

You can use the Banner page to specify Session (exec), Login, and Message-of-the-Day (motd) banners for a security appliance or shared policy.

Note If you use the token `$(hostname)` or `$(domain)` in a banner, it is replaced with the host name or domain name of the security appliance. When you enter the `$(system)` token in a context configuration, the context uses the banner configured in the system configuration.

Spaces in banner text are preserved; however, tabs cannot be entered. Multiple lines in a banner are created by entering a separate line of text for each line you wish to add. Each line is then appended to the end of the existing banner. If the line is empty, a carriage return (CR) is added to the banner.

There is no limit on the length of a banner other than RAM and flash-memory limits. You can only use ASCII characters, including new-line (press the Enter key), which counts as two characters. When accessing the security appliance through Telnet or SSH, the session closes if there is not enough system memory available to process the banner messages, or if a TCP write error occurs when attempting to display the banner messages.

Related Topics

- Chapter 52, "Configuring Server Access Settings on Firewall Devices"

Step 1 To configure banners, access the Banner page:

- (Device view) Select **Platform > Device Admin > Banner** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Banner** from the Policy Types selector. Select an existing policy from the Policies selector, or create a new one.

Step 2 In the **Session (exec) Banner** field, enter the text you want the system to display as a banner before displaying the enable prompt.

Step 3 In the **Login Banner** field, enter the text you want the system to display as a banner before the password login prompt when accessing the security appliance using Telnet.

Step 4 In the **Message-of-the-Day (motd) Banner** field, enter the text you want the system to display as a message-of-the-day banner.

Step 5 To replace a banner, change the contents of the appropriate box.

Step 6 To remove a banner, clear the contents of the appropriate box.

Configuring Boot Image/Configuration Settings

Use the Boot Image/Configuration page to specify which configuration file the security appliance will use at start-up. You can also specify the path to an Adaptive Security Device Manager (ASDM) configuration file.

If you do not specify a boot-image location, the first valid image on internal flash memory will be chosen to launch the system.

Note This page is available only on ASA and PIX 7.0+ devices

Navigation Path

- (Device view) Select **Platform > Device Admin > Boot Image/Configuration** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Boot Image/Configuration** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 48-5 Boot Image/Configuration Page

Element	Description
Boot Config Location	Enter the path to and name of the configuration file to be used when the system is loaded. On an ASA, you can use any of the following syntactical constructs: • **disk0:/**[*path/*]*filename* The value "disk0" represents the internal flash card. You can also use "**flash**" instead of "**disk0**," as they are aliased: – **flash:/**[*path/*]*filename* • **disk1:/**[*path/*]*filename* The value "disk1" represents the external flash card. On a PIX device, you can use only the "flash" syntax; that is: • **flash:/**[*path/*]*filename*
ASDM Image Location	The location and name of the ASDM software image to be used when ASDM sessions are initiated. (You can use ASDM to monitor both ASA and PIX devices.) On a PIX device, as with the Boot Config Location, you are restricted to only the "flash" syntax. On an ASA, as with the Boot Config Location, you can use the "disk0," "flash," or "disk1" constructs. In addition, you can specify an image file on a TFTP server, as follows • **tftp://**[*user*[*:password*]@]*server*[*:port*]/[*path/*]*filename*
Boot Images Table	This table lists any alternate configuration files you have defined; you can specify up to four. The first available image in this list is used if you did not specify a primary file in the Boot Config Location field, or if that file is unavailable. This is a standard Security Manager table; use the up-arrow, down-arrow, Add Row, Edit Row, and Delete Row buttons below the table to manage these entries, as described in Using Tables, page 1-48. The Add Row and Edit Row open the Images Dialog Box, page 48-12, used to add and edit the paths to alternate configuration files. **Note** On an ASA, the first (and only the first) entry in this table can refer to an ASDM configuration file on a TFTP server. If the device cannot reach the TFTP server, it will attempt to load the next image file in the list.

Images Dialog Box

Use the Images dialog box to add or edit a configuration file entry in the Boot Images table on the Boot Image/Configuration page.

Navigation Path

You can access the Images dialog box from the Boot Image/Configuration page. For more information, see Configuring Boot Image/Configuration Settings, page 48-10.

Field Reference

The Images dialog box contains one field, used to define the path to a boot image or configuration file, as follows:

Table 48-6 Images Dialog Box

Element	Description
Image File	Enter the path to and name of the configuration file to add to the ordered Boot Images list. On a PIX device, you can use only the "flash" syntax; that is: • **flash:/**[*path/*]*filename* On an ASA, you can use any of the following syntactical constructs: • **disk0:/**[*path/*]*filename* The value "disk0" represents the internal flash card. You can also use "**flash**" instead of "**disk0**," as they are aliased: – **flash:/**[*path/*]*filename* • **disk1:/**[*path/*]*filename* The value "disk1" represents the external flash card. In addition, on an ASA, you can specify an ASDM image file on a TFTP server, as follows • **tftp://**[*user*[*:password*]@]*server*[*:port*]/[*path/*]*filename* Note that you can specify only one TFTP location, and it must be listed first in the Boot Images table on the Boot Image/Configuration page.

Configuring CLI Prompt

You can use the CLI Prompt page to customize the prompt used by ASA 7.2(1)+ devices during CLI sessions. By default, the prompt shows the hostname of the ASA. In multiple context mode, the prompt also displays the context name. You can display the following items in the CLI prompt:

Note The attributes that are available differ by ASA version:

cluster-unit (ASA 9.1.1+ only)	Displays the cluster unit name. Each unit in a cluster can have a unique name.

context	(Multiple mode only) Displays the name of the current context.
domain	Displays the domain name.
hostname	Displays the hostname.
management-mode (ASA 9.2.1+ only)	Displays the management mode.
priority	Displays the failover priority as pri (primary) or sec (secondary).
state	Displays the traffic-passing state or role of the unit. For failover, the following values appear for the state: • act—Failover is enabled, and the unit is actively passing traffic. • stby— Failover is enabled, and the unit is not passing traffic and is in a standby, failed, or another inactive state. • actNoFailover—Failover is not enabled, and the unit is actively passing traffic. • stbyNoFailover—Failover is not enabled, and the unit is not passing traffic. This condition might occur when there is an interface failure above the threshold on the standby unit. For clustering, the following values are displayed for state: • master • slave For example, in the prompt ciscoasa/cl2/slave, the hostname is ciscoasa, the unit name is cl2, and the state is slave.

Step 1 Access the CLI Prompt page by doing one of the following:

- (Device view) Select **Platform > Device Admin > CLI Prompt** from the Device Policy selector.

  Note For devices in multiple context mode, the CLI Prompt page is only available in the system context. In the admin context, the CLI Prompt page is not available.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > CLI Prompt** from the Policy Types selector. Select an existing policy from the Policies selector, or create a new one.

Step 2 Customize the CLI prompt by doing the following:

- To add an attribute to the prompt, select the attribute in the Available Members list and then click **>>**. The attribute is moved from the Available Members list to the Selected Members list.

  You can add multiple attributes to the prompt. The order in which the attributes are added to the Selected Members list will dictate the order in which they are shown in the CLI prompt.

  Note For ASA 9.1.1+, you can configure up to six attributes for the CLI prompt. For earlier ASA versions, you can only configure up to five attributes.

- To remove an attribute from the prompt, select the attribute in the Selected Members list and then click <<. The attribute is moved from the Selected Members list to the Available Members list.

Setting the Device Clock

Use the Clock page to set the date and time on the selected device.

Note This page is not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

To dynamically set the time using an NTP server, see NTP Page, page 52-21; time derived from an NTP server overrides any time set manually on the Clock page.

Note In multiple-context mode, set the time in the system context only.

Navigation Path

- (Device view) Select **Platform > Device Admin > Clock** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Clock** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 48-7 Clock Page

Element	Description
Device Time Zone	Choose a time zone for the device. These options represent Greenwich Mean Time (GMT) plus or minus a number of hours. **Note** Changing the time zone on the device may drop the connection to any installed Security Services modules (SSMs).
Daylight Savings Time (Summer Time)	Choose a daylight savings time option, and if necessary, specify when and how daylight savings time is applied: **None** – Choose this option to disable automatic correction for daylight savings time. **Set by Date** – Choose this option to specify the date and time when daylight savings time begins and ends for a specific year. If you use this option, you will need to reset the dates every year. **Set by Day** – Choose this option to specify the start and end dates for daylight saving time using the month, week, and day on which daylight savings time begins and ends. This option also lets you set a recurring date range that you do not need to alter yearly.
Set by Date The following three parameters appear in a Start section and an End section. Use the two sets to define when daylight savings time starts and ends.	

Table 48-7 Clock Page (continued)

Element	Description
Date	Enter the date on which daylight savings time begins or ends, in MMM DD YYYY format (for example, Jul 15 2011). You also can click the calendar icon to select the date from a pop-up calendar.
Hour	Choose the hour, from 00 to 23, in which daylight savings time begins or ends.
Minute	Choose the minute, from 00 to 59, at which daylight savings time begins or ends.
Set by Day	
Specify Recurring Time	Check this box to enable the Start and End parameters, which are used to set a recurring start and end for daylight savings time that you do not need to alter yearly.
The following five parameters appear in a Start section and an End section. Use the two sets to define when daylight savings time starts and ends.	
Month	Choose the month in which daylight savings time begins or ends.
Week	Choose the week of the month in which daylight savings time begins or ends. You can select a numerical value that corresponds to the week—1 through 4—or you can specify the first or last week in the month by choosing **first** or **last**. For example, if the day might fall in the partial fifth week, choose last.
Weekday	Choose the day of the week on which daylight savings time begins or ends.
Hour	Choose the hour, from 0 to 23, in which daylight savings time begins or ends.
Minute	Choose the minute, from 00 to 59, at which daylight savings time begins or ends.

Enabling/Disabling FIPS

Beginning with 4.15, Cisco Security Manager provides an option to enable or disable the Federal Information Processing Standards (FIPS) mode on the ASA devices. When you enable FIPS mode with FOM, instead of legacy methods implemented in Cisco SSL versions, the FIPS 140-2 standard compliant cryptographic methods implemented in the FOM is used for signature and verification purposes. This feature is supported only on ASA 9.8.2 or higher devices.

Note To configure FIPS mode on a device, you must reboot the device manually.

Before you enable FIPS, ensure the following are configured on ASA:

1. DH group is set to 14 or ECDH group is set to 19, 20, or 21.
2. The Device Identity Certificate key type is set to RSA and the key size is 2048 or above.

Navigation Path

- (Device view) Select **Platform > Device Admin > FIPS** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > FIPS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 48-8 FIPS Page

Element	Description
FIPS	Select the check box to enable FIPS on the device. This option is available only for ASA 9.8.2 or above. **Note** You must reboot the device after enabling or disabling FIPS for the configuration to take effect.

Configuring Umbrella Global Policy

Beginning with version 4.18, Cisco Security Manager supports configuring Umbrella global policy. Cisco Umbrella Branch is a cloud based security service, which first inspects the DNS traffic, and then examines suspicious HTTP/(S) traffic. The Umbrella connector intercepts DNS packets and redirects the interesting DNS queries to the Umbrella resolver for resolution. Once the DNS response is received, it forwards the response to the host. This feature is supported only on ASA 9.10.1 or higher devices.

After configuring Umbrella service, ensure that the Umbrella DNS policy-map is also configured.

Navigation Path

- (Device view) Select **Platform > Device Admin > Umbrella** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Umbrella** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 48-9 Umbrella Page

Element	Description
Umbrella	Select the check box to apply the global Umbrella configuration for the selected device (ASA 9.10.1 or higher).
Token	The token value of the ASA device on registering with the Umbrella server. Cisco Security Manager throws an error message if this value is less than 64 characters.
Public Key	The public key value of the ASA device on registering with the Umbrella server. This value must be 64 hexa-decimal digits and less than 80 characters, else Cisco Security Manager will display an error message.
EDNS Flow Timeout	The configured EDNS timeout value. The Umbrella time-out for edns-flow must be between <0:0:0> and <1193:0:0>, else Cisco Security Manager will display error message.

Configuring Device Credentials

Use the Credentials page to specify the user credentials Security Manager will use when contacting this device. You can also change the Enable password and the Telnet/SSH password on the device.

This user name-password combination lets you log into the device (in EXEC mode) if you connect to the security appliance using an HTTP, HTTPS, Telnet, or SSH session. You also can specify a separate password specifically for Telnet and SSH sessions. (Further, you can define separate credentials for HTTP/HTTPS connections on the Device Credentials Page, page 3-45 in the Device Properties window.)

The Enable password lets you access privileged EXEC mode after you log in.

Tip The Username, Password, and Enable Password on this page are linked to the Credentials settings in the Device Properties window. When you update these parameters and then deploy the changes to the device, Security Manager uses the existing credentials defined in the Device Properties to log into the device and deploy changes. After successful deployment, the Device Properties credentials are updated to match these settings. For more information about Credentials in Device Properties, see Device Credentials Page, page 3-45.

Navigation Path

- (Device view) Select **Platform > Device Admin > Credentials** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Credentials** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Warning **Since several user accounts may exist on each device, applying a shared Credentials policy to multiple devices will update only the Enable Password on each device; the Username and Password provided in the shared policy are not applied (nor is the Telnet/SSH password). The Enable Password is sufficient for subsequent access to PIX/ASA/FWSM devices, unless external authentication such as AAA or TACACS+ is configured, in which case the Enable Password alone is not sufficient. In this situation, you must manually update the Username, Password and Enable Password on each device that uses external authentication.**

Related Topics

- Configuring User Accounts, page 51-7

Field Reference

Table 48-10 Credentials Page

Element	Description
Username	Enter a user name for logging into the device. The name must be at least four characters; the maximum is 64 characters. Entries are case-sensitive.
Password Confirm	Provide a password for logging into the device (EXEC mode) with the specified Username. This password must be at least three characters; the maximum is 32 characters. Entries are case-sensitive. Re-enter the user password in the Confirm field. **Note** We recommend a password length of at least 8 characters.

Table 48-10 Credentials Page (continued)

Element	Description
Privilege Level	Choose a privilege level for this user; available values are 1 through 15. Level 1 allows EXEC-mode access only, the default level for log-in; level 15 allows Privileged EXEC-mode access; that is, access to the Enable mode. Other levels must be explicitly defined on the device.
Enable password	
Password as encrypted	Select Plain Text or Encrypted.
Password encrypt type	Select MD5 or PBKDF2.
Enable Password Confirm	You can specify an Enable password that lets this user access Privileged EXEC mode after logging in. Entries are case-sensitive. Re-enter the Enable password in the Confirm field. **Note** For Plain Text passwords: – The length of MD5 password should be three to 32 characters. – The length of PBKDF2 password should be 33 to 127 characters. **Note** If you configure user authentication for Enable access, each user has their own password and this password is not used. See Configuring AAA - Authentication Tab, page 48-5 for more information.
Telnet/SSH Password Confirm	You can specify a password that provides access to EXEC mode when connecting to the device via a Telnet or SSH session. This password must be at least three characters; the maximum is 32 characters. Entries are case-sensitive. Re-enter the Telnet/SSH password in the Confirm field. **Note** If you configure user authentication for Telnet or SSH access, each user has their own password and this password is not used. See Configuring AAA - Authentication Tab, page 48-5 for more information.

Managing Mount Points

Use the Mount Points page to make a Common Internet File System (CIFS) or a File Transfer Protocol (FTP) file system accessible to the security appliance.

Note When you create an FTP-type mount point, the FTP server must have a UNIX directory listing style. Microsoft FTP servers have the MS-DOS directory listing style as their default.

The File Mount Point Configuration table lists the configured mount points. The File Mount Point Configuration table is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons. (These are standard buttons, as described in Using Tables, page 1-48.) The Add Row button opens the Add Mount Point Configuration dialog box, and Edit Row opens the Edit Mount Point Configuration dialog box. Other than the titles, the two dialog boxes are identical. For more information, see Add/Edit Mount Point Configuration Dialog Box, page 48-19.

Note This feature is available only on ASA 8.0(2)+ devices. Mount points are supported in router mode only. For ASA versions between 8.0(2) and 9.x, mount points are not supported in multiple-context mode. Mount points are supported in the Admin context on ASA 9.x+ devices in multiple-context, routed mode.

Navigation Path

- (Device view) Select **Platform > Device Admin > Mount Points** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Mount Points** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Add/Edit Mount Point Configuration Dialog Box

Use the Add/Edit Mount Point Configuration dialog box to add or edit a mount point entry in the File Mount Point Configuration table on the Mount Points page. You use mount points to make a Common Internet File System (CIFS) or a File Transfer Protocol (FTP) file system accessible to the security appliance.

Navigation Path

You can access the Add/Edit Mount Point Configuration dialog box from the Mount Points page. For more information, see Managing Mount Points, page 48-18.

Field Reference

Table 48-11 Add/Edit Mount Point Configuration Dialog Box

Element	Description
Enable Mount Point	Whether the file system is mounted or unmounted (available or unavailable).
Connection Type	Select the type of file system to mount: • **cifs**—Specifies that the file system being mounted is CIFS, a file system that provides volume-mounting capabilities for CIFS-shared directories. • **ftp**—Specifies that the file system being mounted is FTP, a Linux kernel module, enhancing the Virtual File System (VFS) with FTP volume-mounting capabilities that allow you to mount FTP-shared directories. **Note** When you create an FTP-type mount point, the FTP server must have a UNIX directory listing style. Microsoft FTP servers have the MS-DOS directory listing style as their default.
Mount Point Name	Specifies the name of the mount point. The mount point name is used when other CLI commands refer to the filesystem already mounted on the security appliance. A maximum of 31 characters are allowed for the mount point name.
Server Name/IP Address	Specifies the predefined name (or the IP address in dotted decimal notation) of the CIFS or FTP file-system server.
Username	Specifies the authorized username for file-system mounting.

Table 48-11 Add/Edit Mount Point Configuration Dialog Box (continued)

Element	Description
Password Confirm	Identifies the authorized password for file-system mounting.
Encrypt Password	Select to indicate that the password supplied is in encrypted format.
Share Name (CIFS only)	Explicitly identifies a specific server share (a folder) by name to access file data within a server.
Domain Name (CIFS only)	For CIFS file systems only, this argument specifies the Windows NT domain name. A maximum of 63 characters is permitted.
Mode (FTP only)	Identifies the FTP transfer mode as either active or passive.
Path (FTP only)	Specifies the directory pathname to the specified FTP file-system server. Question marks and spaces are not allowed in the pathname and will be suppressed.

IP Client

The IP Client page lists the interface name and the IP version. You can use the configured IP Client for integrated routing and bridging support on Firepower 2100 Series devices. The IP Client page has the standard options to Add, Edit and Delete the entries.

Note This feature is available only on ASA 9.8.2+ Firepower 2100 Series single context devices. No multi context support for IP-Client.

Navigation Path

- (Device view) Select **Platform > Device Admin > IP Client** from the Device Policy selector.

Note The menu appears only for a Firepower 2100 Series device.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > IP Client** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Add/Edit IP Client Dialog Box

Use the Add/Edit IP Client dialog box to add or edit a IP Client entry in the IP Client table on the IP Client page. You use IP Client settings to support integrated routing and bridging on Firepower 2100 Series devices.

Navigation Path

You can access the Add/Edit IP Client dialog box from the IP Client page. For more information, see IP Client, page 48-20.

Field Reference

Table 48-12 Add/Edit IP Client Dialog Box

Element	Description
IP Version	The IP address of the device. It can be an IPv4 or an IPv6 address.
Interface	Select the interface pertaining to the Firepower 2100 Series device.

The preview configuration page displays the IP Client configuration with IPv6 suffixed for IPv6 interfaces. For IPv4 interfaces, only the interface name is displayed.

App Agent

Use the App Agent page to configure the App Agent settings. You can specify the heartbeat interval and retry count.

Note App-Agent is available only on Firepower 2100 Series, Firepower 4000 Series, and Firepower 9000 Series devices. In Cisco Security Manager, App-Agent in Firepower 2100 Series device is supported from 9.8.2+; App-Agent in Firepower 4000 Series, and Firepower 9000 Series device is supported from 9.6.2+.

Navigation Path

- (Device view) Select **Platform > Device Admin > App Agent** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > App Agent** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 48-13 App Agent Page

Element	Description
Interval	Enter the app agent heartbeat interval. From ASA 9.6.2 to ASA 9.8.1, App-Agent heartbeat value can be between 300 to 6000 ms. For ASA 9.8.2+ devices, App-Agent heartbeat interval value can be between 100 to 6000 ms. **Note** Cisco Security Manager will display an error message if you do not enter values in multiples of 100.
Retry Count	Enter the retry count between 3 to 10.
Save	Click to save the configuration.

Configuring Device Access Settings on Firewall Devices

The Device Access section, located under the Device Admin folder in the Policy selector, contains pages for defining access to firewall devices.

This chapter contains the following topics:

- Configuring Console Timeout, page 49-1
- HTTP Page, page 49-2
- Configuring ICMP, page 49-4
- Configuring Management Access, page 49-6
- Configuring Secure Shell Access, page 49-7
- Configuring SSL - Basic and Advanced tabs, page 49-8
- Reference Identities, page 49-13
- Configuring SNMP, page 49-14
- Telnet Page, page 49-29

Configuring Console Timeout

Use the Console page to specify a timeout value for inactive console sessions. When the time limit you specify is reached, the console session is closed.

In the **Console Timeout** field, enter the number of minutes a console session can remain idle before the device closes it. Valid values are 0 to 60 minutes. To prevent a console session from timing out, enter 0.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Console** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > Console** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 49, "Configuring Device Access Settings on Firewall Devices"

HTTP Page

Use the table on the HTTP page to manage the interfaces configured to access the HTTP server on a device, as well as HTTP redirect to HTTPS on those interfaces. You also can enable or disable the HTTP server on the device from this page. Administrative access by the specific device manager requires HTTPS access.

Note To redirect HTTP, the interface requires an access list that permits HTTP. Otherwise, the interface cannot listen to port 80, or to any other port that you configure for HTTP.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > HTTP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > HTTP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 49-1 HTTP Page

Element	Description
HTTP Interface table	Use the Add Row, Edit Row, and Delete Row buttons below this table to manage device interfaces on which HTTP-to-HTTPS redirect is configured. Add Row and Edit Row open the HTTP Configuration Dialog Box, page 49-3.
Fetch user name from certificate settings	Select this option to set the rules for extracting a username from the certificate. Enter the following: • Enable HTTP username from certificate—Check this box to get the HTTP username from the certificate for authentication. • Pre-fill user name—Check the Pre-fill Username checkbox to enable the use of this name for authentication. When enabled, this username, along with the password entered by the user, are used for authentication. Choose one of the following options: **Note** This feature is supported only in devices running ASA software version 9.4(1) or higher. • Use the Entire DN as the username—Select this option if you want to use the entire DN as the username. This option is disabled by default.

Table 49-1 HTTP Page (continued)

Element	Description
Fetch user name from certificate settings (contd..)	• Specify individual DN fields as the username—Choose from the Primary DN Field and Secondary DN Field drop-down values to specify which attributes and additional attributes to use to derive the username. This option is enabled by default. – C—Country: the two-letter country abbreviation which conforms to the ISO 3166 country abbreviations. – CN—Common Name: the name of a person, system, or other entity. Not available as a secondary attribute. – DNQ—Domain Name Qualifier. – EA—Email address. – GENQ—Generational qualifier. – GN—Given name. – I—Initials. – L—Locality: the city or town where the organization is located. – N—Name. – O—Organization: the name of the company, institution, agency, association, or other entity. – OU—Organizational Unit: the subgroup within the organization (O). – SER—Serial number. – SN—Surname. – SP—State/Province: the state or province where the organization is located. – T—Title. – UID—User Identifier. – UPN—User Principal Name. • Use LUA Script generated by ASDM—Choose this option if you want to use the LUA script that is generated by ASDM. This option is disabled by default.
Enable HTTP Server	Enables or disables the HTTP server on the device. When enabled, you can specify a communications **Port** for the server. The Port range is 1 to 65535; the default is 443.

HTTP Configuration Dialog Box

Use the HTTP Configuration dialog box to add or edit a host or network that will be allowed to access the HTTP server on the device via a specific interface; you also can enable and disable HTTP redirect.

Navigation Path

You can access the HTTP Configuration dialog box from the HTTP Page, page 49-2.

Field Reference

Table 49-2 HTTP Configuration Dialog Box

Element	Description
Interface Name	Enter or Select the interface on which access to the HTTP server on the device is allowed. **Note** Beginning with Cisco Security Manager version 4.17, you can configure BVI interface for HTTP on ASA 9.9.2 devices and above. However, in multi-context, "Transparent" mode security context only supports BVI interface.
IP Address/Netmask	Enter the IP address and netmask, separated by a forward slash ("/") of the host or network that is permitted to establish an HTTP connection with the device. Alternately, you can click Select to select a Networks/Hosts object. **Note** Beginning with version 4.13, Cisco Security Manager supports policies—Groups, Hosts, Address Range, and Network for IPv6 devices.
Enable Authentication Certificate	Select this option to require user certificate authentication in order to establish an HTTP connection. On ASA and PIX 8.0(2)+ devices, you can specify the authentication **Port**.
Certificate Maps	Select the Certificate Map name that you configured in Remote Access VPN > certificate to Connection Profile Maps > Rules. For more information, see Map Rule Dialog Box (Upper Table), page 31-39. None is selected by default. This feature is available beginning with Security Manager version 4.12 for ASA 9.6(2) or later devices. This option is supported for single, multi, routed and transparent contexts for ASA devices.
Redirect port	The port on which the security appliance listens for HTTP requests, which it then redirects to HTTPS. To disable HTTP redirect, ensure that this field is blank.

Configuring ICMP

Use the table on the ICMP page to manage Internet Control Message Protocol (ICMP) rules, which specify the addresses of all hosts or networks that are allowed or denied ICMP access to specific interfaces on the security device.

Note Starting from ASA 8.2(1) ICMP IPv6 was supported in the transparent firewall mode.

The ICMP rules control ICMP traffic that terminates on any device interface. If no ICMP control list is configured, the device accepts all ICMP traffic that terminates at any interface, including the outside interface. However, by default, the device does not respond to ICMP echo requests directed to a broadcast address.

It is recommended that permission is always granted for the ICMP Unreachable message (type 3). Denying ICMP Unreachable messages disables ICMP Path MTU discovery, which can halt IPsec and PPTP traffic. See RFC 1195 and RFC 1435 for details about Path MTU Discovery.

If an ICMP control list is configured, the device uses a first match to the ICMP traffic, followed by an implicit deny all. That is, if the first matched entry is a permit entry, the processing of the ICMP packet continues. If the first matched entry is a deny entry, or an entry is not matched, the device discards the ICMP packet and generates a syslog message. If an ICMP control list is not configured, a permit rule is assumed in all cases.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > ICMP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > ICMP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Note ICMP IPv6 support is not available for PIX and FWSM devices.

Field Reference

Table 49-3 ICMP Page

Element	Description
ICMP Rules Table	Use the Add Row, Edit Row, and Delete Row buttons below this table to manage ICMP rules. Add Row opens the Add ICMP dialog box, while Edit Row opens the Edit ICMP dialog box. See Add and Edit ICMP Dialog Boxes, page 49-5 for information about these dialog boxes.
ICMP Unreachable Parameters	
Rate Limit	For ICMP traffic that terminates at an interface on this device, the maximum number of ICMP Unreachable messages the device can transmit per second. This value can be between 1 and 100 messages per second; the default is 1 message per second.
Burst Size	The burst size for ICMP Unreachable messages; this can be a value between 1 and 10. **Note** This parameter is not currently used by the system, so you can choose any value.

Add and Edit ICMP Dialog Boxes

Use the Add ICMP dialog box to add an ICMP rule, which specifies a host/network that is allowed or denied the specified ICMP access on the specified device interface.

Note The Edit ICMP dialog box is virtually identical to the Add ICMP dialog box, and is used to modify existing ICMP rules. The following descriptions apply to both dialog boxes.

Navigation Path

You can access the Add or Edit ICMP dialog boxes from the Configuring ICMP, page 49-4.

Note While adding an ICMP policy, make sure that the network and service is of the same type i.e. IPv6 networks support IPv6 services.

Field Reference

Table 49-4 Add and ICMP Dialog Boxes

Element	Description
Action	Whether this rule permits or denies the selected ICMP Service message from the specified Network on the specified Interface. Choose: • **Permit** – ICMP messages from the specified networks/hosts are allowed to the specified interface. • **Deny** – ICMP messages from the specified networks/hosts to the specified interface are dropped.
ICMP Service	Enter or Select the specific ICMP service message to which the rule applies.
Interface	Enter or Select the device interface to which these ICMP messages are directed.
Network	Enter a host name, IPv4 or IPv6 address, or Select a Networks/Hosts object, to define the specified ICMP message source.

Configuring Management Access

Use the Management Access page to enable or disable access on a high-security interface so you can perform management functions on the device. You can enable this feature on an internal interface to allow management functions to be performed on the interface over an IPsec VPN tunnel. You can enable the Management Access feature on only one interface at a time.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Management Access** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > Management Access** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Enabling and Disabling Management Access

In the **Management Access Interface** field, enter the name of the device interface that is to permit management access connections. You can click Select to select the interface from a list of interface objects.

You can enable the Management Access feature on only one interface at a time.

Clear the Management Access Interface field to disable management access.

Configuring Secure Shell Access

Use the Secure Shell page to configure rules that permit administrative access to a security device using the SSH protocol. The rules restrict SSH access to a specific IP address and netmask. Any SSH connection attempts that comply with these rules must then be authenticated by an AAA server or Telnet password.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Secure Shell** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > Secure Shell** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 49-5 Secure Shell Page

Element	Description
SSH Version	Specify the SSH version(s) accepted by the device: choose **1**, **2**, or **1 and 2**. By default, SSH Version 1 and SSH Version 2 connections are accepted.
Timeout	Enter the number of minutes, 1 to 60, the Secure Shell session can remain idle before the device closes it. The default value is 5 minutes.
Allowed Hosts table	Use the Add Row, Edit Row, and Delete Row buttons below this table to manage the hosts allowed to connect to the security device via SSH. Add Row opens the Add Host dialog box, while Edit Row opens the Edit Host dialog box. See Add and Edit SSH Host Dialog Boxes, page 49-8 for information about these dialog boxes.
Enable Secure Copy	Check this box to enable the secure copy (SCP) server on the security appliance. This allows the appliance to function as an SCP server for transferring files from/to the device. Only clients that are allowed to access the security appliance using SSH can establish a secure copy connection. This implementation of the secure copy server has the following limitations: • The server can accept and terminate connections for secure copy, but cannot initiate them. • The server does not have directory support. The lack of directory support limits remote client access to the security appliance internal files. • The server does not support banners. • The server does not support wildcards. • The security appliance license must have the VPN-3DES-AES feature to support SSH version 2 connections.

Add and Edit SSH Host Dialog Boxes

Use the Add Host dialog box to add an SSH access rule.

Note The Edit Host dialog box is virtually identical to the Add Host dialog box, and is used to modify existing SSH access rules. The following descriptions apply to both dialog boxes.

Navigation Path

You can access the Add and Edit Host dialog boxes from the Configuring Secure Shell Access, page 49-7.

Field Reference

Table 49-6 Add and Edit Host Dialog Boxes

Element	Description
Interface	Enter or Select the name of the device interface on which SSH connections are permitted. **Note** Beginning with Cisco Security Manager version 4.17, you can configure BVI interface for SSH connections on ASA 9.9.2 devices and above. However, in multi-context, "Transparent" mode security context only supports BVI interface.
IP Addresses	Enter the name or IP address for each host or network that is permitted to establish an SSH connection with the security device on the specified interface; use commas to separate multiple entries. You also can click Select to select Networks/Hosts objects from a list. **Note** Beginning with version 4.13, Cisco Security Manager supports policies—Groups, Hosts, Address Range, and Network for IPv6 devices.

Configuring SSL - Basic and Advanced tabs

Beginning from version 4.8, Security Manager provides enhanced security features using Secure Sockets Layer (SSL).

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SSL** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > SSL** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 49-7 SSL Page Basic Tab

Element	Description
Certificate Authentication	
FCA Timeout	Enter a value in the range of 1 to 120. **Note** FCA Timeout is applicable for devices running the ASA software version 9.1(2) or higher.
Interface	Use the Add Row, Edit Row, and Delete Row buttons below the Interface table to manage the interfaces and their port numbers allowed to connect to the security device via SSL. Add Row opens the Add Interface and Port dialog box, while Edit Row opens the Edit Interface and Port dialog box. You can select the interface from the available entries in the Interface Selector dialog box. Enter a value in the range of 1 to 65535 for the port number.
Client Version SSL/TLS Protocol Version	The Client Version is the SSL/TLS protocol version to use when the device acts as a client. Select any one of the following: • Any—Select this keyword to transmit SSLV3 ClientHellos and negotiate SSLV3 or greater. This is the default keyword. • SSLV3—Enter this keyword to transmit SSLv3 ClientHellos and negotiate SSLV3 or greater. • TLSV1—Enter this keyword to transmit TLSv1 ClientHellos and negotiate TLSV1 or greater. • TLSV1.1—Enter this keyword to transmit TLSV1.1 ClientHellos and negotiate TLSV1.1 or greater. • TLSV1.2—Enter this keyword to transmit TLSV1.2 ClientHellos and negotiate TLSV1.2 or greater. **Note** TLSV1.1 and TLSV1.2 protocol versions are applicable for devices running the ASA software version 9.3(2) or higher.

Table 49-7 SSL Page Basic Tab (continued)

Element	Description
Server Version SSL/TLS Protocol Version	The Server Version is the minimum SSL/TLS protocol version to use when the device acts as a server. Select any one of the following: • Any—Select this keyword to accept SSLV2 ClientHellos and negotiate the highest common version. This is the default keyword. • SSLV3—Enter this keyword to accept SSLV2 ClientHellos and negotiate SSLV3 or greater. • SSLV3-Only—Enter this keyword to accept SSLV2 ClientHellos and negotiate SSLV3 or greater. • TLSV1—Enter this keyword to accept SSLV2 ClientHellos and negotiate TLSV1 or greater. • TLSV1-Only—Enter this keyword to accept SSLV2 ClientHellos and negotiate TLSV1 or greater. • TLSV1.1—Enter this keyword to accept SSLV2 ClientHellos and negotiate TLSV1.1 or greater. • TLSV1.2—Enter this keyword to accept SSLV2 ClientHellos and negotiate TLSV1.2 or greater. **NOTES:** • The **Any** keyword is the default for both Server Version and Client Version and means that the device will negotiate the highest common supported version of TLS. • TLSV1.1 and TLSV1.2 protocol versions are applicable for devices running the ASA software version 9.3(2) or higher. • SSLV3-Only and TLSV1-Only protocol versions are applicable for devices running the ASA software version older than 9.3(2).

Table 49-8 SSL Page Advanced Tab

Element	Description
Advanced SSL Settings for devices running the ASA software version older than 9.3(2)	
Encryption	Choose the encryption algorithms from the available list. To add an encryption algorithm, select the item in the Available Members list and then click >>. The item is moved from the Available Members list to the Selected Members list. You can add multiple encryption algorithms. The available encryption algorithms are as follows: • 3DES-SHA1 • AES128-SHA1 • AES256-SHA1 • DES-SHA1 • RC4-MD5 • RC4-SHA1 • NULL-SHA1 • DHE-AES128-SHA1 • DHE-AES256-SHA1 To remove an encryption algorithm, select the item in the Selected Members list and then click <<. The item is moved from the Selected Members list to the Available Members list. Click Save to save your settings.
Advanced SSL Settings for devices running the ASA software version 9.3(2) or higher	
SSL Cipher	Use the Add Row, Edit Row, and Delete Row buttons below the SSL Cipher table to manage the SSL cipher version and level. On the Add Cipher dialog select a combination of the version and level.
Version	Select one of the following versions: • DEFAULT • DTLSV1 • DTLSV1.2 • SSLV3 • TLSV1 • TLSV1.1 • TLSV1.2 **Note** The DEFAULT keyword is used to configure outbound connections when the device is acting as a client and establishing a connection to a server. All the other keywords are used when the device is acting as a server and accepting connections from a client. **Note** The SSLV3 version has been deprecated from ASA version 9.4(1). Therefore, beginning with version 4.9, Security Manager performs a validation to check if SSLV3 option has been configured for any ASA devices running the version 9.4(1) or higher.

Table 49-8 SSL Page Advanced Tab (continued)

Element	Description
Level	Select one of the following versions: • ALL - It includes all ciphers including NULL-SHA. • LOW - It includes all ciphers except NULL-SHA. • MEDIUM - It includes all ciphers except NULL-SHA:DES-CBC-SHA:RC4-SHA:RC4-MD5. • FIPS - It includes all FIPS-compliant ciphers (that is, not NULL-SHA:DES-CBC-SHA:RC4-MD5:RC4-SHA:DES-CBC3-SHA) • HIGH - It includes only AES-256 with SHA-2 ciphers, so it only applies to TLSV1.2.
Custom String	Use the CUSTOM keyword for Security Manager to exercise full control over the cipher suite using OpenSSL cipher definition strings. **Note** Beginning with version 4.9, Security Manager provides support for the following new TLSV1.2 ciphers for devices running the ASA software version 9.4(1) or higher. • ECDHE_RSA_AES128_SHA256 • ECDHE_RSA_AES256_SHA384 • ECDHE_ECDSA_AES128_SHA256 • ECDHE_ECDSA_AES256_SHA384 **Note** Beginning with version 4.16, Security Manager provides support for the following new TLSV1.2 ciphers in addition to the above mentioned ciphers for devices running the ASA software version 9.4(1) or higher. • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-GCM-SHA384 • AES256-GCM-SHA384 • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • AES128-GCM-SHA256
ECDH Configuration	Select from one of the options (19,20,21,none) in the ECDH Group. This feature is available from Security Manager version 4.9 onwards for ASA devices version 9.4(1) or higher.

Note Due to import regulations in some countries the Oracle implementation provides a default cryptographic jurisdiction policy file that limits the strength of cryptographic algorithms. If stronger algorithms need to be configured or are already configured on the device (for example, AES with 256-bit keys, DH group with 5,14,24), follow these steps:

1. Download the Java 7 unlimited strength cryptography policy .jar files from http://www.oracle.com. Cisco recommends to search for the following on the Oracle website:

   Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files Java 7

(Click the download button to download the files by accepting the license agreement.)

2. Replace local_policy.jar and US_export_policy.jar on your Security Manager server in the folder CSCOpx\MDC\vms\jre\lib\security.
3. Restart your Security Manager server.

Reference Identities

Beginning with version 4.12, Security Manager enables you to configure Reference Identity policy objects for Secure Syslog Server connections on devices running the ASA software version 9.6(2) or later. This object enables support for Common Criteria requirements.
Reference identities are configured as one or more identifiers to be compared to the presented identifiers in the server certificate. Identifiers are specific instances of the four identifier types specified in RFC 6125.

Add/Edit Reference Identity Dialog Box

Use the Add/Edit Reference Identity Dialog Box to create a new Reference Identity policy object or to edit existing policy objects.

Navigation Path

Select **Manage > Policy Objects**, then select **Reference Identity** from the Object Type Selector. Right-click inside the work area, then select **New Object** or click the **+** button to add a new object, or right-click a row, then select **Edit Object.**

Field Reference

Table 49-9 Add/Edit Reference Identity Dialog Box

Element	Description
Name	Name of the Reference Identity policy object. Note that each Reference Identifier can have multi line values.
Description	Description of the Reference Identity policy object.
Common Name ID	A Relative Distinguished Name (RDN) in a certificate subject field that contains only one attribute-type-and-value pair of type Common Name (CN), where the value matches the overall form of a domain name. The CN value can be free text. A CN-ID reference identifier does not identify an application service.
Domain Name ID	A subjectAltName entry of type dNSName. This is a DNS domain name. A DNS-ID reference identifier does not identify an application service.
Service Name ID	A subjectAltName entry of type otherName whose name form is SRVName as defined in RFC 4985. A SRV-ID identifier may contain both a domain name and an application service type. For example, a SRV-ID of “_imaps.example.net” would be split into a DNS domain name portion of “example.net” and an application service type portion of “imaps.”

Table 49-9 Add/Edit Reference Identity Dialog Box (continued)

Element	Description
Uniform Resource Identifier ID	A subjectAltName entry of type uniformResourceIdentifier whose value includes both (i) a "scheme" and (ii) a "host" component (or its equivalent) that matches the "reg-name" rule specified in RFC 3986. A URI-ID identifier must contain the DNS domain name, not the IP address, and not just the hostname. For example, a URI-ID of "sip:voice.example.edu" would be split into a DNS domain name portion of "voice.example.edu" and an application service type of "sip."
Category	(Optional) Select a category between CAT-A and CAT-J.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the Edit button to create, edit, and view the overrides. The Overrides field indicates the number of devices that have overrides for this object.

Note A reference identity is created when configuring one with a previously unused name. Once a reference identity has been created, the four identifier types and their associated values can be added or deleted from the reference identity. The reference identifiers MAY contain information identifying the application service and MUST contain information identifying the DNS domain name.

Configuring SNMP

Simple Network Management Protocol (SNMP) defines a standard way for network management stations running on PCs or workstations to monitor the health and status of many types of devices, including switches, routers, and security appliances. You can use the SNMP page to configure a firewall device for monitoring by SNMP management stations.

The Simple Network Management Protocol (SNMP) enables monitoring of network devices from a central location. Cisco security appliances support network monitoring using SNMP versions 1, 2c, and 3, as well as traps and SNMP read access; SNMP write access is not supported.

You can configure a security appliance to send "traps" (event notifications) to a network management station (NMS), or you can use the NMS to browse the management information bases (MIBs) on the security appliance. Use CiscoWorks for Windows or any other SNMP MIB-II-compliant browser to receive SNMP traps and browse a MIB.

The security appliance has an SNMP agent that notifies designated management stations if specified events occur, for example, when a link in the network goes up or down. The notification includes an SNMP system object ID (OID), identifying the device to the management stations. The security appliance SNMP agent also replies when a management station asks for information.

SNMP MIBs and OIDs

An SNMP trap reports significant events occurring on a network device, most often errors or failures. SNMP traps are defined in Management Information Bases (MIBs), which can be either standard or enterprise-specific.

Standard traps and MIBs are created by the Internet Engineering Task Force (IETF) and documented in various RFCs. Standard traps are compiled into the security appliance software. If needed, you can also download RFCs, standard MIBS, and standard traps from the IETF website: http://www.ietf.org/.

For Cisco MIB files and OIDs, refer to: http://www.cisco.com/public/sw-center/netmgmt/cmtk/mibs.shtml. OIDs may be downloaded from this FTP site: ftp://ftp.cisco.com/pub/mibs/oid/oid.tar.gz.

This section contains the following topics:

- SNMP Terminology, page 49-15
- SNMP Version 3, page 49-15
- SNMP Page, page 49-17

SNMP Terminology

Here are definitions for some common SNMP terms:

- **Agent** – The SNMP server running on the security appliance. The agent responds to requests for information and action from the management station. The agent also controls access to its management information base (MIB), the collection of data objects that can be viewed or changed by the SNMP manager.
- **Management stations** – The PCs or workstations set up to monitor SNMP events and manage devices such as the security appliance. Management stations can also receive messages about events which require attention, such as hardware failures.
- **MIBs** – The agent maintains standardized data structures called Management Information Bases (MIBs), used to collect information, such as packet, connection and error counters, and buffer usage and failover status. A number of MIBs are defined for specific products, and for the common protocols and hardware standards used by most network devices. SNMP management stations can browse MIBs, or request only specific fields. In some applications, MIB data can be modified for administrative purposes.
- **OID** – The SNMP standard assigns a system object ID (OID) so that a management station can uniquely identify network devices with SNMP agents, and indicate to users the source of information monitored and displayed.
- **Traps** – Specified events that generate a message from the SNMP agent to the management station. Events include alarm conditions such as linkup, linkdown, coldstart, warmstart, authentication, or syslog events.

SNMP Version 3

SNMP Version 3 provides security enhancements that are not available in SNMP Version 1 or SNMP Version 2c. SNMP Versions 1 and 2c transmit data between the SNMP server and SNMP agent in clear text. SNMP Version 3 adds authentication and privacy options to secure protocol operations. In addition, this version controls access to the SNMP agent and MIB objects through the User-based Security Model

(USM) and View-based Access Control Model (VACM). The ASA and ASASM also support the creation of SNMP groups and users, as well as hosts, which is required to enable transport authentication and encryption for secure SNMP communications.

Note SNMP Version 3 is supported on ASA devices running 8.2(1) or later and on ASASM devices running 8.5(1) or later.

Security Models

For configuration purposes, the authentication and privacy options are grouped together into security models. Security models apply to users and groups, which are divided into the following three types:

- NoAuth—No Authentication and No Privacy, which means that no security is applied to messages.
- Auth—Authentication but No Privacy, which means that messages are authenticated.
- Priv—Authentication and Privacy, which means that messages are authenticated and encrypted.

SNMP Groups

An SNMP group is an access control policy to which users can be added. Each SNMP group is configured with a security model, and is associated with an SNMP view. A user within an SNMP group must match the security model of the SNMP group. These parameters specify what type of authentication and privacy a user within an SNMP group uses. Each SNMP group name and security model pair must be unique.

SNMP Users

SNMP users have a specified username, a group to which the user belongs, authentication password, encryption password, and authentication and encryption algorithms to use. The authentication algorithm options are MD5 and SHA. The encryption algorithm options are DES, 3DES, and AES (which is available in 128, 192, and 256 versions). When you create a user, you must associate it with an SNMP group. The user then inherits the security model of the group.

SNMP Hosts

An SNMP host is an IP address to which SNMP notifications and traps are sent. To configure SNMP Version 3 hosts, along with the target IP address, you must configure a username, because traps are only sent to a configured user. SNMP target IP addresses and target parameter names must be unique on the ASA and ASA Services Module. Each SNMP host can have only one username associated with it. To receive SNMP traps, configure the SNMP NMS, and make sure that you configure the user credentials on the NMS to match the credentials for the ASA and ASASM.

Implementation Differences Between the ASA, ASA Services Module, and the Cisco IOS Software

The SNMP Version 3 implementation in the ASA and ASASM differs from the SNMP Version 3 implementation in the Cisco IOS software in the following ways:

- The local-engine and remote-engine IDs are not configurable. The local engine ID is generated when the ASA or ASASM starts or when a context is created.
- No support exists for view-based access control, which results in unrestricted MIB browsing.
- Support is restricted to the following MIBs: USM, VACM, FRAMEWORK, and TARGET.
- You must create users and groups with the correct security model.
- You must remove users, groups, and hosts in the correct sequence.

- Use of the **snmp-server host** command creates an ASA or ASASM rule to allow incoming SNMP traffic.

SNMP Page

Use the SNMP page to configure the security appliance for monitoring by Simple Network Management Protocol (SNMP) management stations.

Note For SNMP Version 3, configuration must occur in the following order: group, user, host.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > SNMP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Host Access Entry Dialog Box, page 49-22
- Add/Edit SNMP Host Group Entry Dialog Box, page 49-23
- Add/Edit SNMP Group Entry Dialog Box, page 49-24
- Add/Edit SNMP User Entry Dialog Box, page 49-25
- Add/Edit SNMP User List Entry Dialog Box, page 49-27

Field Reference

Table 49-10 SNMP Page

Element	Description
Enable SNMP Servers	When this option is selected, the security device provides SNMP information on the specified interface(s). You can deselect this option to disable SNMP monitoring while retaining the configuration information.
Read Community String Confirm	Enter the password used by a SNMP management station when sending requests to this device. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. The security device uses the password to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces are not permitted. Repeat the password in the **Confirm** field to ensure it was entered correctly.

Table 49-10 SNMP Page (continued)

Element	Description
System Administrator Name	Enter the name of the device administrator or other contact person. This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
Location	Describe the location of this security device (for example, ***Building 42, Sector 54***). This string is case-sensitive and can be up to 127 characters. Spaces are accepted, but multiple spaces are shortened to a single space.
Port (PIX 7.x, ASA and FWSM 3.x only)	Specify the port on which incoming requests will be accepted. The default is 161.
Configure SNMP Traps	Click this button to configure SNMP traps in the SNMP Trap Configuration Dialog Box, page 49-19.
SNMP Engine ID	Shows the ID of the SNMP engine configured on the device. Click **Get SNMP Engine ID** to retrieve the engine ID from the device.
SNMP Hosts table	This table lists the SNMP management stations that can access the security appliance. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, as described in Using Tables, page 1-48. The Add Row and Edit Row buttons open the Add/Edit SNMP Host Access Entry Dialog Box, page 49-22, used to add and edit management station host entries. **Note** For ASA devices running 9.1(5) or later, you can configure up to 129 SNMP hosts. For other devices and earlier ASA versions, you can only configure up to 32 SNMP hosts.
SNMP Host Group table	Beginning with version 4.12, Security Manager enables you to add and edit the Host Group entries for SNMP Users. See Add/Edit SNMP Host Group Entry Dialog Box, page 49-23 for more information.
SNMPv3 Configuration	
SNMP Group tab	Lists the SNMP groups that have been configured. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, as described in Using Tables, page 1-48. The Add Row and Edit Row buttons open the Add/Edit SNMP Group Entry Dialog Box, page 49-24, used to add and edit SNMP groups.
SNMP User tab	Lists the SNMP users that have been configured. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, as described in Using Tables, page 1-48. The Add Row and Edit Row buttons open the Add/Edit SNMP User Entry Dialog Box, page 49-25, used to add and edit SNMP users.
SNMP User List tab	Beginning with version 4.12, Security Manager enables you to add a user list containing multiple SNMP users. See Add/Edit SNMP User List Entry Dialog Box, page 49-27 for more information.

SNMP Trap Configuration Dialog Box

Use the SNMP Trap Configuration dialog box to configure SNMP traps (event notifications) for the selected security device.

Traps are different than browsing; they are unsolicited "comments" from the managed device to the management station for certain events, such as ***linkup***, ***linkdown***, and ***syslog event generated***.

An SNMP object ID (OID) for the device appears in SNMP event traps sent from the device. The SNMP service running on a security device performs two functions:

- Replies to SNMP requests from management stations.
- Sends traps to management stations or other devices that are registered to receive them from the security appliance.

Cisco security devices support three types of traps:

- firewall
- generic
- syslog

In the SNMP Trap Configuration dialog box, available traps are presented on four tabbed panels: Standard, Entity MIB, Resource, and Other.

Navigation Path

You can access the SNMP Trap Configuration dialog box from the SNMP Page, page 49-17.

Related Topics

- Configuring SNMP, page 49-14
- Add/Edit SNMP Host Access Entry Dialog Box, page 49-22

Field Reference

Table 49-11 SNMP Trap Configuration Dialog Box

Element	Description
Enable All SNMP Traps	Check this box to quickly select all traps on all four tabbed panels.
Enable Syslog Traps	Check this box to enable transmission of trap-related syslog messages. The severity level for syslog messages trapped is set on the Logging Filters Page, page 54-12.
Select the desired event-notification traps on the following four tabbed panels. Note that only the traps applicable to the selected device are displayed in the dialog box.	

Table 49-11 SNMP Trap Configuration Dialog Box (continued)

Element	Description
Standard	• **Authentication** – Unauthorized SNMP access. This authentication failure occurs for packets with an incorrect community string. • **Link Up** – One of the device's communication links has become available (it has "come up"), as indicated in the notification. • **Link Down** – One of the device's communication links has failed, as indicated in the notification. • **Cold Start** – The device is reinitializing itself such that its configuration or the protocol entity implementation may be altered. • **Warm Start** – The device is reinitializing itself such that its configuration and the protocol entity implementation is unaltered.
Entity MIB	• **Field Replaceable Unit Insert** – A Field Replaceable Unit (FRU) has been inserted, as indicated. (FRUs include assemblies such as power supplies, fans, processor modules, interface modules, etc.) • **Field Replaceable Unit Delete** – A Field Replaceable Unit (FRU) has been removed, as indicated in the notification. • **Configuration Change** – There has been a hardware change, as indicated in the notification. • **Fan Failure** – A device cooling fan has failed, as indicated in the notification. • **CPU Temperature** – Temperature of the central processing unit has reached the configured limit. • **Power-Supply Failure** – A device power supply has failed, as indicated in the notification. • **Redundancy Switchover** – Switchover occurred for redundant component, as indicated in the notification. • **Alarm Asserted** – The condition described by the alarm exists. • **Alarm Cleared** – The condition described by the alarm does not exist.
Resource	• **Connection Limit Reached** – This trap indicates that a connection attempt was rejected because the configured connections limit has been reached. • **Resource Limit Reached** – This notification is generated when the configured resource limit is reached, as described in the notification. • **Resource Rate Limit Reached** – This notification is generated when the configured resource rate limit is reached, as described in the notification

Table 49-11 SNMP Trap Configuration Dialog Box (continued)

Element	Description
Other	• **IKEv2 Start** – Internet Key Exchange version 2 (IKEv2) exchange initiated. • **IKEv2 Stop** – Internet Key Exchange version 2 (IKEv2) exchange terminated. • **Memory Threshold** – Available free memory has fallen below configured threshold, as indicated in the notification. • **ASA CPU Rising Threshold** – This notification is triggered when utilization of CPU resources exceeds the specified Percentage for a specified Period of time: **Percentage** – Enter the desired upper limit of CPU resource usage as a percentage of total available. Valid values range from 10 to 94; default is 70%. **Period** – Specify the length of time, in minutes, that the specified Percentage can be exceeded before notification is triggered. Valid values range from 1 to 60. • **Interface Threshold** – This notification is triggered when utilization of a physical interface exceeds the specified Percentage of total bandwidth: **Percentage** – Enter the desired upper limit on interface usage as a percentage of total available bandwidth. Valid values range from 30 to 99; default is 70%. • **IPSec Start** – IPsec has started, as indicated in the notification. • **IPSec Stop** – IPsec has stopped, as indicated in the notification. • **Remote Access Session Threshold Exceeded** – The number of remote access sessions has reach the defined limit, as indicated in the notification. • **NAT Packet Discard** – This notification is generated when IP packets are discarded by the NAT function. Available Network Address Translation addresses or ports have fallen below configured threshold. • **CPU Rising Threshold** – This notification is triggered when utilization of CPU resources exceeds the specified Percentage for a specified Period of time: **Percentage** – Enter the desired upper limit of CPU resource usage as a percentage of total available. Valid values range from 10 to 100; default is 70%. **Period** – Specify the length of time, in seconds, that the specified Percentage can be exceeded before notification is triggered. Valid values range from 60 to 3600.

Add/Edit SNMP Host Access Entry Dialog Box

Use the Add/Edit SNMP Host Access Entry dialog box to add and edit entries in the SNMP Hosts table on the SNMP page. These entries represent SNMP management stations allowed to access the security device.

For ASA devices running any software version between 9.1(5) and 9.3(2), you can configure 129 SNMP hosts. For ASA devices running the software version lower than 9.1(5) you can configure only 32 SNMP hosts.

Beginning with version 4.9, Security Manager enables you to configure up to 4096 SNMP hosts for ASA devices running the software version 9.4(1) or higher. However, only 129 of this number can be for traps. You cannot configure more than 129 trap configured SNMP hosts.

Navigation Path

You can access the Add/Edit SNMP Host Access Entry dialog box from the SNMP Page, page 49-17.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Group Entry Dialog Box, page 49-24
- Add/Edit SNMP User Entry Dialog Box, page 49-25

Field Reference

Table 49-12 Add/Edit SNMP Host Access Entry Dialog Box

Element	Description
Interface Name	Enter or Select the interface on which this SNMP management station contacts the device.
IP Address	Enter the IP address, or Select a Networks/Hosts object, representing the SNMP management station. **Note** Beginning with Cisco Security Manager version 4.17, IPv6 Address for SNMP policy is supported on ASA 9.9.2 devices and above.
UDP Port	(Optional) Enter a UDP port for requests from the SNMP host. You can use this field to override the global value specified on the SNMP page.
Community String Confirm	Enter the password used by the SNMP management station when sending requests to the security device. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. Thus, the password is used to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces are not permitted. Repeat the password in the **Confirm** field.
SNMP Version	Choose the version of SNMP used by this management station: **1**, **2c**, or **3**.
SNMP User Name	If SNMP version 3 is selected, select the SNMP user. For information on SNMP users, see Add/Edit SNMP User Entry Dialog Box, page 49-25.

Table 49-12 Add/Edit SNMP Host Access Entry Dialog Box (continued)

Element	Description
Server Poll/Trap Specification	Specify the type(s) of communication with this management station: poll only, trap only, or both trap and poll. Check either or both: • Poll – The security device waits for periodic requests from the management station. • Trap – The device sends trap events when they occur.

Add/Edit SNMP Host Group Entry Dialog Box

Beginning with Security Manager version 4.12, you can use the Add/Edit SNMP Host Group Entry dialog box to add and edit entries in the SNMP Host Group table on the SNMP page. These entries represent SNMP management stations allowed to access the security device.

For ASA devices running any software version between 9.1(5) and 9.4, you can configure 129 SNMP hosts. For ASA devices running the software version lower than 9.1(5) you can configure only 32 SNMP hosts.

Beginning with version 4.9, Security Manager enables you to configure up to 4096 SNMP hosts for ASA devices running the software version 9.4(1) or higher. However, only 129 of this number can be for traps. You cannot configure more than 129 trap configured SNMP hosts.

Note If you edit a used Address Range or Network object in the Networks/Host Policy Object Manager after adding or editing SNMP Host or Host Group entries in the Add/ Edit SNMP Host Group entry page, Cisco Security Manager will not validate for the total number of SNMP traps. Thus, if the trap entries exceed 129, it will result in a deployment failure.

Navigation Path

You can access the Add/Edit SNMP Host Access Entry dialog box from the SNMP Page, page 49-17.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Group Entry Dialog Box, page 49-24
- Add/Edit SNMP User Entry Dialog Box, page 49-25

Field Reference

Table 49-13 Add/Edit SNMP Host Access Entry Dialog Box

Element	Description
Interface Name	Enter or Select the interface on which this SNMP management station contacts the device.
IP Address	Enter the IP address, or Select a Networks/Hosts object, representing the SNMP management station.
UDP Port	(Optional) Enter a UDP port for requests from the SNMP host. You can use this field to override the global value specified on the SNMP page.

Table 49-13 Add/Edit SNMP Host Access Entry Dialog Box (continued)

Element	Description
Community String Confirm	Enter the password used by the SNMP management station when sending requests to the security device. The SNMP community string is a shared secret among the SNMP management stations and the network nodes being managed. Thus, the password is used to determine if the incoming SNMP request is valid. The password is a case-sensitive alphanumeric string of up to 32 characters; spaces are not permitted. Repeat the password in the **Confirm** field.
SNMP Version	Choose the version of SNMP used by this management station: **1**, **2c**, or **3**.
Server Poll/Trap Specification	Specify the type(s) of communication with this management station: poll only, trap only, or both trap and poll. Check either or both: • Poll – The security device waits for periodic requests from the management station. • Trap – The device sends trap events when they occur. **Note** You cannot enable both traps and polling for the same SNMP Host Group. If you need to enable this, Cisco recommends that you use the **snmp-server host** command for the relevant hosts.

Add/Edit SNMP Group Entry Dialog Box

Use the Add/Edit SNMP Group Entry dialog box to add and edit entries in the SNMP Groups table on the SNMP page. An SNMP group is an access control policy to which users can be added. Each SNMP group is configured with a security model, and is associated with an SNMP view. A user within an SNMP group must match the security model of the SNMP group. These parameters specify what type of authentication and privacy a user within an SNMP group uses. Each SNMP group name and security model pair must be unique.

For configuration purposes, the authentication and privacy options are grouped together into security models. Security models apply to users and groups, which are divided into the following three types:

- NoAuth—No Authentication and No Privacy, which means that no security is applied to messages.
- Auth—Authentication but No Privacy, which means that messages are authenticated.
- Priv—Authentication and Privacy, which means that messages are authenticated and encrypted.

Notes

- Before a group can be deleted, you must ensure that all users associated with that group are deleted. If any hosts are associated with a user that needs to be deleted, you must delete those hosts before you can delete the user.
- If users have been configured to belong to a particular group with a certain security model, you must do the following to change the security level of that group:
  a. Remove all host entries associated with any users belonging to the group.
  b. Remove the users from the group.
  c. Deploy the changes to the device.
  d. Change the group security level.

e. Add users that belong to the group.

f. Add hosts belonging to the users that were added for the group.

g. Deploy the changes to the device.

Navigation Path

You can access the Add/Edit SNMP Group Entry dialog box from the SNMP Page, page 49-17.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Host Access Entry Dialog Box, page 49-22
- Add/Edit SNMP User Entry Dialog Box, page 49-25

Field Reference

Table 49-14 Add/Edit SNMP Group Entry Dialog Box

Element	Description
Group Name	Enter the name of the SNMP group. Group names must be 32 characters or less.
Security Level	Specify the security level for the group: • NoAuth—No Authentication and No Privacy, which means that no security is applied to messages. • Auth—Authentication but No Privacy, which means that messages are authenticated. • Priv—Authentication and Privacy, which means that messages are authenticated and encrypted.

Add/Edit SNMP User Entry Dialog Box

Use the Add/Edit SNMP User Entry dialog box to add a user to an SNMP group or to edit entries in the SNMP User table on the SNMP page. SNMP users inherit the security model of the group to which they are assigned.

Notes

- After a user has been created, you cannot change the group to which the user belongs.
- Before a user can be deleted, you must ensure that no hosts are configured that are associated with that username.

Navigation Path

You can access the Add/Edit SNMP User Entry dialog box from the SNMP Page, page 49-17.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Host Access Entry Dialog Box, page 49-22

- Add/Edit SNMP Group Entry Dialog Box, page 49-24

Field Reference

Table 49-15 Add/Edit SNMP User Entry Dialog Box

Element	Description
Group Name	Select the SNMP group to which this user belongs. For information on SNMP groups, see Add/Edit SNMP Group Entry Dialog Box, page 49-24.
Security Level	Shows the security level for the selected group: • NoAuth—No Authentication and No Privacy, which means that no security is applied to messages. • Auth—Authentication but No Privacy, which means that messages are authenticated. • Priv—Authentication and Privacy, which means that messages are authenticated and encrypted.
User Name	Enter the name of the SNMP user. Usernames must be 32 characters or less and must be unique for the SNMP server group selected.
Engine ID (SNMP version v3 only)	The SNMP EngineID identifier used for authentication in v3. You can enter comma separated multiple Engine IDs. The Engine ID identifier must be valid, and each Engine Id must be within the range of 1 to 257 characters. **Note** • If you configure EngineID for an SNMP user with MD5 algorithm, the EngineID must be a valid one. If the EngineID is not valid, the preview config would fail with an error "failed to generate raw config". For example, the preview config fails if the EngineID entered is 111. • For an SNMP group with a security level of NoAuth, do not provide an EngineID identifier because on deployment, the ASA will ignore this engine ID and take the default local engine ID. • The following dynamic behaviors of the device cannot be handled in Security Manager: – If you upgrade a failover ASA device from version 8.x or 9.x to version 9.6(2), the device will automatically create multiple SNMP User commands for multiple SNMP Engine IDs. You must copy the Engine ID by retrieving it from the device into this Engine ID text box. For information about retrieving Engine ID from the device see SNMP Page, page 49-17. – If you add or remove an ASA device to or from a failover configuration, you must manually enter the Engine ID because the ASA device automatically removes or creates new SNMP User commands for the existing Engine IDs.

Table 49-15 Add/Edit SNMP User Entry Dialog Box (continued)

Element	Description
Encrypt Password Type	Specify the type of password you want to use: Clear Text or Encrypted. If the password type is Clear Text, Security Manager will encrypt the password when deploying to the device. If the password type is Encrypted, Security Manager will directly deploy the encrypted password. Security Manager will never directly deploy the clear text password to device.
Auth Algorithm Type	Specify the type of authentication you want to use: MD5 or SHA.
Authentication Password Confirm	Enter the password to use for authentication. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as *xx:xx:xx...*, where *xx* are hexadecimal values. **Note** The length of the password will depend on the authentication algorithm selected. For all passwords, the length must be 256 characters or less. If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.
Encryption Type	Specify the type of encryption you want to use: AES128, AES192, AES256, 3DES, DES. **Note** To use AES or 3DES encryption, you must have the appropriate license installed on the device.
Encryption Password Confirm	Enter the password to use for encryption. If you selected Encrypted as the Encrypt Password Type, the password must be formatted as *xx:xx:xx...*, where *xx* are hexadecimal values. For encrypted passwords, the length of the password depends on the encryption type selected. The password sizes are as follows (where each *xx* is one octal): • AES 128 requires 16 octals • AES 192 requires 24 octals • AES 256 requires 32 octals • 3DES requires 32 octals • DES can be any size **Note** For all passwords, the length must be 256 characters or less. If you selected Clear Text as the Encrypt Password Type, repeat the password in the **Confirm** field.

Add/Edit SNMP User List Entry Dialog Box

Beginning with version 4.12, Security Manager enables you to use the Add/Edit SNMP User List Entry dialog box to add a user list containing multiple SNMP users.

Notes

- You cannot delete a user list if the list is used by a particular host group.

- You cannot delete an SNMP user if that user is referred to in a particular user list..

Navigation Path

You can access the Add/Edit SNMP User List Entry dialog box from the SNMP Page, page 49-17.

Field Reference

Table 49-16 Add/Edit SNMP User List Entry Dialog Box

Element	Description
User List Name	Enter the name of the User List. User List names must be 1-33 characters in length.
User Names	Select the user names from the drop-down list.

Related Topics

- Configuring SNMP, page 49-14
- SNMP Trap Configuration Dialog Box, page 49-19
- Add/Edit SNMP Host Access Entry Dialog Box, page 49-22
- Add/Edit SNMP Group Entry Dialog Box, page 49-24

Telnet Page

Use the Telnet page to configure rules that permit only specific hosts or networks to connect to the firewall device using the Telnet protocol.

The rules restrict administrative Telnet access through a firewall device interface to a specific IP address and netmask. Connection attempts that comply with the rules must then be authenticated by a preconfigured AAA server or the Telnet password. You can monitor Telnet sessions using Monitoring > Telnet Sessions.

Note Only five Telnet sessions can be active at the same time in single-context mode. In multiple-context mode on ASAs, there can be only five Telnet sessions active per context, 100 Telnet sessions active per blade. With resource class, the administrator can further tune this parameter.

Related Topics

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Telnet** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Device Access > Telnet** from the Policy Type selector. Right-click **Telnet** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Telnet Configuration Dialog Box, page 49-29

Field Reference

Table 49-17 Telnet Page

Element	Description
Timeout	Number of minutes a Telnet session can remain idle before the firewall device closes it. Values can range from 1 to 1440 minutes.
Telnet Access Table	
Interface	Interface that receives Telnet packets from the client.
IP Addresses	The IP address and network mask of each host or network that can access the Telnet console through the specified interface.

Telnet Configuration Dialog Box

Use the Telnet Configuration dialog box to configure Telnet options for an interface.

Navigation Path

You can access the Telnet Configuration dialog box from the Telnet Page, page 49-29.

Field Reference

Table 49-18 Telnet Configuration Dialog Box

Element	Description
Interface Name	Enter or Select an interface that can receive Telnet packets from a client. **Note** Beginning with Cisco Security Manager version 4.17, you can configure BVI interface for Telnet on ASA 9.9.2 devices and above. However, in multi-context, "Transparent" mode security context only supports BVI interface.
IP Addresses/Netmask	Enter or Select the IP address and netmask, separated by a "/", of each host or network permitted to access the firewall device's Telnet console through the specified interface. Use commas to separate multiple entries. **Note** To limit access to a single IP address, use 255.255.255.255 or 32 as the netmask. Do not use the subnetwork mask of the internal network. **Note** Beginning with version 4.13, Cisco Security Manager supports policies—Groups, Hosts, Address Range, and Network for IPv6 devices.

Configuring Failover

The Failover page provides access to failover settings for the selected security appliance. The available settings and the overall appearance of the Failover page may change slightly, depending upon the type of device selected, its mode of operation (routed or transparent), and its context mode (single or multiple).

In other words, how you configure failover depends upon both the operating mode and the security context of the security appliance.

Please note the following caveats when assigning an interface as a failover link:

- You can define the interface in the Add/Edit Interface dialog box, but do not configure it. In particular, **do not specify an interface Name**, as this parameter disqualifies the interface from being used as the failover link. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for more information.
- IPv6 addresses are not supported for failover links.
- On an ASA 5505, an interface assigned as the backup for another interface cannot be used as a failover link (although no checking is performed to prevent this).
- Do not assign a PPPoE-enabled interface as a failover link. PPPoE and Failover should not be configured on the same device interface (although no checking is performed to prevent this).
- A failover interface cannot use the same IP address as another interface, especially the Management IP address (although no checking is performed to prevent this).

Note also that after you assign an interface as a failover link, the interface is listed on the Interfaces page, but you cannot edit or delete the interface from that page. The only exception is if you set a physical interface to be the stateful failover link—you can configure its speed and duplex.

This chapter contains the following topics:

- Understanding Failover, page 50-1
- Basic Failover Configuration, page 50-5
- Additional Steps for an Active/Standby Failover Configuration, page 50-9
- Failover Policies, page 50-10

Understanding Failover

Failover lets you configure two identical security appliances such that one will take over firewall operations if the other fails. Using a pair of security appliances, you can provide high system availability without operator intervention.

The linked security appliances communicate failover information over a dedicated link. This failover link can be either a LAN-based connection or, on PIX security appliances, a dedicated serial failover cable. The following information is communicated over the failover link:

- Current failover state (active or standby)
- "Hello" messages (also called "keep-alives")
- Network link status
- MAC address exchange
- Configuration replication
- Per-connection state information, in the case of Stateful failover

Caution All information sent over the failover link is sent in clear text unless you secure the communication with a failover key. If the security appliance is used to terminate VPN tunnels, this information includes any user names, passwords, and preshared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing failover communications with a failover key, particularly if you are using the security appliance to terminate VPN tunnels.

Cisco security appliances support two types of failover:

- **Active/Standby** – The *active* security appliance inspects all network traffic, while the *standby* security appliance remains idle until a failure occurs on the active appliance. Changes to the configuration of the active security appliance are transmitted over the failover link to the standby security appliance.

  When failover occurs, the standby security appliance becomes the active unit, and it assumes the IP and MAC addresses of the previously active unit. Because other devices on the network do not see any changes in the IP or MAC addresses, ARP entries do not change or time-out anywhere on the network.

  Active/Standby failover is available to security appliances operating in single- or multiple-context mode. In single-context mode, only Active/Standby failover is available, and all failover configuration is by means of the Failover page.

  Note When using Active/Standby failover, you must make all configuration changes on the active unit. The active unit automatically replicates the changes to the standby unit. The standby unit should not be imported or added to the Security Manager device list.

  Also, you must manually copy the authentication certificate from the active device to the standby device. See Additional Steps for an Active/Standby Failover Configuration, page 50-9 for additional information.

- **Active/Active** – Both security appliances inspect network traffic by alternating their roles—such that one is active and one is standby—on a per context basis. This means Active/Active failover is available only on security appliances operating in multiple-context mode.

  However, Active/Active failover is not required in multiple-context mode. That is, on a device operating in multiple-context mode, you can configure Active/Standby or Active/Active failover. In either case, you provide system-level failover settings in the system context, and context-level failover settings in the individual security contexts.

  See Active/Active Failover, page 50-3 for additional information about this topic.

In addition, failover can be stateless or stateful:

- **Stateless** – Also referred to as "regular" failover. With stateless failover, all active connections are dropped when failover occurs. Clients need to re-establish connections when the new active unit takes over.
- **Stateful** – The active unit in the failover pair continually passes per-connection state information to the standby unit. When failover occurs, the same connection information is available on the new active unit. Supported end-user applications are not required to reconnect to maintain the current communication session.

  See Stateful Failover, page 50-4 for more information.

Related Topics

- Chapter 50, "Configuring Failover"
- Basic Failover Configuration, page 50-5
- Failover Policies, page 50-10

Active/Active Failover

Active/Active failover is available only on security appliances operating in multiple-context mode. In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device.

The active and standby roles are assigned over the complete set of security contexts, more or less arbitrarily.

To enable Active/Active failover on the security appliance, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments on the unit that will have failover group 1 in the active state. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default.

As in Active/Standby failover, each unit in an Active/Active failover pair is given a primary or secondary designation. Unlike Active/Standby failover, this designation does not indicate which unit is active when both units start simultaneously. Each failover group in the configuration is given a primary or secondary role preference. This preference determines the unit on which the contexts in the failover group appear in the active state when both units start simultaneously. You can have both failover groups be in the active state on a single unit in the pair, with the other unit containing the failover groups in the standby state. However, a more typical configuration is to assign each failover group a different role preference to make each one active on a different unit, balancing the traffic across the devices.

Note To reliably manage security contexts in Active/Active failover mode, Cisco Security Manager requires an IP address for the management interface of each context so that it can communicate directly with the active security context of a failover pair.

Initial configuration synchronization occurs when one or both units start. This synchronization occurs as follows:

- When both units start simultaneously, the configuration is synchronized from the primary unit to the secondary unit.
- When one unit starts while the other unit is already active, the unit that is starting up receives the configuration from the already active unit.

After both units are running, commands are replicated from one unit to the other as follows:

- Commands entered within a security context are replicated from the unit on which the security context is in the active state to the peer unit.

  Note A context is considered in the active state on a unit if the failover group to which it belongs is in the active state on that unit.

- Commands entered in the system execution space are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.
- Commands entered in the admin context are replicated from the unit on which failover group 1 is in the active state to the unit on which failover group 1 is in the standby state.

Failure to enter the commands on the appropriate unit for command replication to occur will cause the configurations to be out of synchronization. Those changes may be lost the next time the initial configuration synchronization occurs.

Note When bootstrapping the peer devices in an Active/Active Failover configuration, the bootstrap configurations are only applied to the system contexts of the respective failover peer devices.

In an Active/Active failover configuration, failover occurs on a failover group basis, not a system basis. For example, if you designate both failover groups as active on the primary unit, and failover group 1 fails, failover group 2 remains active on the primary unit, while failover group 1 becomes active on the secondary unit.

Note When configuring Active/Active failover, make sure that the combined traffic for both units is within the capacity of each unit.

Stateful Failover

Note Stateful failover is not supported on the ASA 5505 appliance.

When stateful failover is enabled, the active unit in the failover pair continually updates the current connection-state information on the standby unit. When failover occurs, supported end-user applications are not required to reconnect to maintain the current communication session.

Note The IP and MAC addresses for the state and LAN failover links do not change at failover.

To employ stateful failover, you must configure a link to pass all state information to the standby unit. If you are using a LAN failover connection rather than the serial failover interface (which is available only on the PIX platform), you can use the same interface for the state link and the failover link. However, we recommend that you use a dedicated interface for passing state information to the standby unit.

The following information is passed to the standby unit when stateful failover is enabled:

- NAT translation table

- TCP connection table (except for HTTP), including the timeout connection
- HTTP connection states (if HTTP replication is enabled)
- H.323, SIP and MGCP UDP media connections
- The system clock
- The ISAKMP and IPsec SA table

The following information is not copied to the standby unit when stateful failover is enabled:

- HTTP connection table (unless HTTP replication is enabled)
- The user authentication (UAUTH) table
- The ARP table
- Routing tables

Basic Failover Configuration

The following steps describe basic failover configuration. Please note the following caveats when assigning an interface as a failover link:

- You can define the interface in the Add/Edit Interface dialog box, but do not configure it. In particular, **do not specify an interface Name**, as this parameter disqualifies the interface from being used as the failover link.
- On an ASA 5505, an interface assigned as the backup for another interface cannot be used as a failover link (although no checking is performed to prevent this).
- Do not assign a PPPoE-enabled interface as a failover link. PPPoE and Failover should not be configured on the same device interface (although no checking is performed to prevent this).
- A Failover interface cannot use the same IP address as another interface, especially the Management IP address (although no checking is performed to prevent this).

Note When you save a failover configuration, it is applied to both the security appliance and the failover peer.

Before You Begin

Licenses installed on the device must allow failover configurations. On ASA 5505 and 5510 devices, this failover license is an optional license. You must install the failover license outside of Security Manager, using ASDM or the device CLI, and ensure that the **License Supports Failover** option is selected in the General page of the device properties (right-click the device and select **Device Properties**). If the license is installed when you add the device to the inventory, or you install the license and then rediscover device policies, Security Manager can identify the license and set this option appropriately.

If the option is selected and the license is not in fact installed, you will see deployment failures. If the option is not selected, Security Manager will not deploy the failover policy to the device even if you configure the policy.

Related Topics

- Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26
- Understanding Failover, page 50-1
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

- Failover Policies, page 50-10

Step 1 Ensure Device View is your present application view; if necessary, click the **Device View** button on the toolbar.

Note For more information on using the Device View to configure device policies, see Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30.

Step 2 Select the appliance you want to configure.

Step 3 Expand the **Platform** entry in the Device Policy selector, then expand **Device Admin** and select **Failover**.

The Failover page is displayed.

Step 4 (PIX only) Choose the **Failover Method**: **Serial Cable** or **LAN Based**. If you choose Serial Cable, the LAN Failover settings are disabled; be sure the cable connecting the two devices is in place.

Step 5 Select **Enable Failover** to enable failover on this appliance.

Step 6 (Optional) Click the Settings button to open the Settings dialog box for the selected device. The contents of the Settings dialog box depend on the type of device, and whether it is operating in single or multiple mode—some options may not be available. Refer to the following sections:

- Settings Dialog Box, page 50-21 (ASA/PIX 7+)
- Advanced Settings Dialog Box, page 50-16 (FWSM)

Step 7 Click the **Bootstrap** button to open the Bootstrap configuration for LAN failover dialog box, which provides bootstrap configurations that can be applied to the primary and secondary devices in a LAN failover configuration. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.

Step 8 (Multiple-context devices only) In the Configuration section, select the failover mode: **Active/Active** or **Active/Standby**.

Step 9 (Optional) Follow these steps to configure an interface for **LAN Failover** communications between the two devices:

a. Assign a device **Interface** for LAN-based communications, and then press the Tab key on your keyboard to update the page.

On PIX and ASA devices, this drop-down list displays the interfaces defined on the device. You can type in a port ID (e.g., *gigabitethernet1*), or you can choose the port if you have already defined the interface.

On an FWSM, the Interface list is not populated with VLAN IDs; you must enter the numeric ID of the VLAN you wish to use.

Note In both cases, this cannot be a Named interface, nor can the interface be configured for PPPoE.

b. Provide a **Logical Name** for this failover interface.

c. Enter the **Active IP** address for failover communications.

d. Enter a **Standby IP** address for failover communications. The Standby IP address is used on the security appliance that is currently the standby unit.

e. Enter the **Subnet Mask** for both IP addresses. Both must be on the same subnet.

Step 10 (Optional) Follow these steps to enable and configure an interface for **Stateful Failover** communications between the two devices:

a. Assign a device **Interface** for update communications, and then press the Tab key on your keyboard to update the page.

You can type in a port ID (e.g., *gigabitethernet1*), or you can choose the port if you have already defined the interface; note that this cannot be a Named interface.

Note On an FWSM, this is a **VLAN** interface.

b. Provide a **Logical Name** for this interface.

c. Enter the **Active IP** address for connection updates.

d. Enter a **Standby IP** address for update communications.

e. Enter the **Subnet Mask** for both IP addresses. Both must be on the same subnet.

f. Select **Enable HTTP Replication** to preserve HTTP connection information.

Connection information is communicated to the standby unit for all TCP protocols except HTTP, because HTTP connections are generally short-lived. Select this option to maintain HTTP connections during failover.

Step 11 Provide a communications-encryption key: enter a **Shared Key** and then repeat it in the **Confirm** field. Be sure to enter the same key on both devices. (Not available on FWSM versions prior to 3.1)

The Shared Key can be any arbitrary string of up to 63 alphanumeric characters. If **HEX** is checked, the Shared Key is an arbitrary string of exactly 32 hexadecimal characters. (The HEX option is available only on PIX/ASA version 7.0.5 and later, and FWSM versions 3.1.3 and later.)

Note This step is optional, but we strongly recommend encrypting failover communications.

Step 12 To specify a failover reconnect timeout value for asymmetrically routed sessions, enter a length of time in the **Timeout** field, in the form hh:mm:ss (the minutes and seconds values are optional). If the field is blank (the default), or contains a zero, reconnections are prevented. Setting this value to -1 disables the timeout, allowing connections to reconnect after any amount of time.

Step 13 (Optional) You can configure Bidirectional Forwarding Detection (BFD), to communicate with a failover pair and this can be used to monitor the health of the failover unit. Create or select a **BFD template** from the **Health-Check Monitoring** section.

Note This is applicable only for Firepower failover devices running ASA 9.7.1 and above.

Tip BFD failover commands are supported only in the Active/Standby mode.
In a multi-context device, BFD failover commands are supported only in the system context.
BFD Failover commands are not supported in the transparent mode.

Step 14 (FWSM only) – Configured interfaces are listed in the Interface Configuration table. To edit the failover configuration for a listed interface, select it and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23.

Adding A Security Context to Failover Group 2

To add a new security context to an existing failover group 2, you must save the new context configuration to a deployment file and then manually add it to the appropriate device. Otherwise, until the first successful deployment, Security Manager will attempt to communicate with the new context through the device's Admin context. This will fail since group 2 cannot be reached through the Admin context (unless both group 1 and 2 are active on the same device).

The following steps outline creating a new security context and adding it to failover group 2.

1. Create the new security context.

   Be sure to define: context Name, Configuration URL, assign an Interface, choose Failover Group 2, and provide a Management IP Address. See Managing Security Contexts, page 59-7 for more information.

2. Save and submit these changes.

3. Provide the following context-configuration information, saving each change as you go:

   - On the Credentials page of the Device Properties window for the new context, provide Username and Password. See Viewing or Changing Device Properties, page 3-41 for additional information.
   - On the context's Interfaces page, edit the assigned interface, providing a Name, IP address and Subnet Mask. See Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26 for additional information.
   - On the context's Failover Page (ASA/PIX 7.0+), page 50-17, edit the interface configuration to provide a Standby IP Address.
   - On the HTTP Page, page 49-2, check Enable HTTP Server and then define HTTP access.
   - On the Credentials page, provide the Username and Password to be used when contacting the context. See Configuring Device Credentials, page 48-17 for additional information.

4. Choose **Deploy** from the Configuration Manager's File menu. Submit your changes, and then in the Deploy Saved Changes dialog box, be sure only this new context is selected, and then click Edit Deploy method. In the Edit Deploy Method dialog box, change the Method to File and then specify the Destination and a file name. Click OK to close the Edit Deploy Method dialog box, and then click Deploy the Deploy Saved Changes dialog box.

   The context configuration is saved to the specified file. See Deploying to a File, page 8-11 for more information about this step.

5. After uploading the configuration file to the device, use the CLI to enable HTTP access for the context. For example:

```
ciscoasa/group2(config-if)#  int g3/0
ciscoasa/group2(config-if)# nameif man
ciscoasa/group2(config-if)# security-level 100
ciscoasa/group2(config-if)# ip add 203.0.113.176 255.255.254.0 st 203.0.113.177
ciscoasa/group2(config-if)# exit
ciscoasa/group2(config)# http serv ena
ciscoasa/group2(config)# http 0.0.0.0 0.0.0.0 man
ciscoasa/group2(config)# username cisco pass cisco
ciscoasa/group2(config)#wr
```

Following this process, any new changes to the context can be successfully deployed to the context with Security Manager (attempts to reach the context will not go through the Admin context's management IP address).

Alternative

Another approach to this issue is to add the new context to failover group 1 first, and then perform the configuration via Security Manager. However, in order to then move this context to failover group 2, both groups (1 and 2) must be active on the same device. Otherwise, this error will be reported:

```
"join-failover-group 2
ERROR: Command requires failover-group 2 and 1 to be in the same state or no nameif comand for all interfaces
in this context"
```

Additional Steps for an Active/Standby Failover Configuration

Cisco Security Manager lets you authenticate a PIX/ASA/FWSM device by validating the certificate installed on the device. When configuring firewalls in an active/standby failover configuration, you must manually copy the certificate from the active device to the standby device so that Security Manager can communicate with the standby device after a failover occurs.

The following procedures describe how to export or display the identity certificate, CA certificate, and keys for a security appliances in your network using ASDM, and then import that information onto a standby device using ASDM.

- Exporting the Certificate to a File or PKCS12 data, page 50-9
- Importing the Certificate onto the Standby Device, page 50-9

Exporting the Certificate to a File or PKCS12 data

To export a trustpoint configuration, follow these steps using ASDM:

Step 1 Go to Configuration > Features > Device Administration > Certificate > Trustpoint > Export.

Step 2 Fill in the Trustpoint Name, Encryption Passphrase, and Confirm Passphrase fields. For information on these fields, click Help.

Step 3 Select a method for exporting the trustpoint configuration.

- Export to a File—Type the filename or browse for the file.
- Display the trustpoint configuration in PKCS12 format—Display the entire trustpoint configuration in a text box and then copy it for importing. For more information, click Help.

Step 4 Click **Export**.

Importing the Certificate onto the Standby Device

To import a trustpoint configuration, follow these steps using ASDM:

Step 1 Go to Configuration > Features > Device Administration > Certificate > Trustpoint > Import.

Step 2 Fill in the Trustpoint Name, Decryption Passphrase, and Confirm Passphrase fields. For information on these fields, click Help. The decryption passphrase is the same as the encryption passphrase used when the trustpoint configuration was exported.

Step 3 Select a method for importing the trustpoint configuration.

- Import from a File—Type the filename or browse for the file.
- Enter the trustpoint configuration in PKCS12 format—Paste the entire trustpoint configuration from the exported source into a text box. For more information, click Help.

Failover Policies

This section lists the pages that describe configuring failover on various types of security appliances; the pages are organized by device type.

PIX 6.x Firewalls

- Failover Page (PIX 6.3), page 50-10
  - Edit Failover Interface Configuration Dialog Box (PIX 6.3), page 50-12
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Firewall Services Modules

- Failover Page (FWSM), page 50-13
  - Advanced Settings Dialog Box, page 50-16
  - Add/Edit Interface MAC Address Dialog Box, page 50-23
  - Edit Failover Interface Configuration Dialog Box, page 50-23
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Adaptive Security Appliances and PIX 7.0 Firewalls

- Failover Page (ASA/PIX 7.0+), page 50-17
  - Settings Dialog Box, page 50-21
  - Edit Failover Group Dialog Box, page 50-25
  - Edit Failover Interface Configuration Dialog Box, page 50-23
  - Add/Edit Interface MAC Address Dialog Box, page 50-23
  - Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Failover Page (PIX 6.3)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure failover settings for a PIX 6.3.x Firewall.

Navigation Path

Select a PIX 6.3.x device in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10

Field Reference

Table 50-1 Failover Page (PIX 6.3)

Element	Description
Failover	
Failover Method	Choose the type of failover link: **Serial Cable** or **LAN Based**. If you choose Serial Cable, ensure the physical cable is connected to both devices.
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key type, flash memory, and RAM. On PIX devices with LAN Based chosen as the Failover Method, you must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.
Failover Poll Time	Specify the amount of time between hello messages among units. Values can range from 3 to 15 seconds; default is 15.
LAN-Based Failover These fields are available when LAN Based is the chosen Failover Method.	
Interface	Choose the interface to be used for LAN-based failover. If "Not Selected" is chosen, LAN-based failover is disabled.
Shared Key Confirm	Used to encrypt communications between the primary and standby devices. Value can be any alphanumeric string. Re-enter the Shared Key in the Confirm field.
Stateful Failover (Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
Interface	Choose the interface to be used for Stateful Failover. If "Not Selected" is chosen, Stateful Failover is disabled. **Note** You must choose a fast LAN link from the list (for example, 100full, 1000full, or 1000sxfull).
Enable HTTP Replication	When selected, active HTTP sessions are copied to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Interface Configuration The table lists all available named interfaces. To define a Standby IP address and Active and Standby MAC addresses for an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box (PIX 6.3), page 50-12.	

Edit Failover Interface Configuration Dialog Box (PIX 6.3)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Edit Failover Interface Configuration dialog box to configure failover interfaces for the selected PIX 6.3.x device.

Note The failover interface cannot be configured for PPPoE.

Navigation Path

You can access the Edit Failover Interface Configuration dialog box from the Interface Configuration table on the Failover Page (PIX 6.3), page 50-10.

Related Topics

- Failover Policies, page 50-10

Field Reference

Table 50-2 Edit Failover Interface Configuration Dialog Box (PIX 6.3)

Element	Description
Interface	The name of the interface; read-only.
Active IP Address	Displays the IP address of the active interface. This address is used by the standby device to communicate with the active device. The address must be on the same network as the system IP address. The active IP address of this interface; read-only. This address is used by the standby device to communicate with the active device. This field is blank if an IP address has not been assigned to the interface. **Tip** You can use this IP address with the ping tool to check the status of the active device.
Netmask	The subnet mask for the active IP address; read-only. This field is blank if an IP address has not been assigned to the interface.
Standby IP Address	Specify the IP address of the corresponding interface on the standby failover unit. This address is used by the active device to communicate with the standby device. The address must be on the same network as the system IP address. This field does not appear if an IP address has not been assigned to the interface. **Tip** You can use this IP address with the ping tool to check the status of the standby device.
Failover MAC Addresses	
These parameters let you define virtual MAC addresses for a physical interface that is configured for failover; these addresses are optional.	

Table 50-2 Edit Failover Interface Configuration Dialog Box (PIX 6.3) (continued)

Element	Description
Active MAC Address	Specify a MAC address for the active interface in hexadecimal format (for example, 0123.4567.89ab).
Standby MAC Address	Specify a MAC address for the standby interface in hexadecimal format (for example, 0123.4567.89ab).

Failover Page (FWSM)

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure basic failover settings for the selected Firewall Services Module (FWSM).

Navigation Path

To access this feature, select a FWSM in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9
- Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Field Reference

Table 50-3 Failover Page (FWSM)

Element	Description
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key, flash memory, and RAM. You must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Settings button	Click to display the Advanced Settings Dialog Box, page 50-16, used to define when failover should occur.
Configuration	
This section is presented only for FWSM 3.1.1+ devices operating in multiple-context mode.	

Table 50-3 Failover Page (FWSM) (continued)

Element	Description
Active/Active	In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device. To enable Active/Active failover on the device, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments on the unit that will have failover group 1 in the active state. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default. See Add/Edit Security Context Dialog Box (FWSM), page 59-8 for information about assigning a context to a failover group.
Active/Standby	In an Active/Standby configuration, the active security appliance handles all network traffic passing through the failover pair. The standby security appliance does not handle network traffic until a failure occurs on the active security appliance. Whenever the configuration of the active security appliance changes, it sends configuration information over the failover link to the standby security appliance. When a failover occurs, the standby security appliance becomes the active unit. It assumes the IP and MAC addresses of the previously active unit. Because the other devices on the network do not see any changes in the IP or MAC addresses, ARP entries do not change or time out.
LAN Failover	
VLAN	Enter the numeric ID of the VLAN interface you are using for the failover link; for example, 11. This list is not automatically populated with VLAN IDs—you must highlight “Not Selected” and type the desired VLAN ID number; press your keyboard’s Tab key to activate the related fields. When configured for failover, the interface is directly connected to the standby device.
Logical Name	Enter a logical name for the failover VLAN interface.
Active IP Address	Specify the active IP address for this interface.
Standby IP Address	Specify a standby IP address for this interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The Standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.
Subnet Mask	Enter the Subnet Netmask for the Active and Standby IP addresses.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.

Table 50-3 Failover Page (FWSM) (continued)

Element	Description
Stateful Failover	
(Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
VLAN	Enter the numeric ID of the VLAN interface you are using for the failover link; for example, 12. This list is not automatically populated with VLAN IDs—you must highlight "Not Selected" and type the desired VLAN ID number; press your keyboard's Tab key to activate the related fields. When configured for failover, the interface is directly connected to the standby device.
Logical Name	Enter a logical name for the Stateful failover VLAN interface.
Active IP Address	Specify the active IP address for this interface.
Standby IP Address	Specify a standby IP address for this interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The Standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.
Subnet Mask	Enter the Subnet Netmask for the Active and Standby IP addresses.
Enable HTTP Replication	When selected, allows stateful failover to copy active HTTP sessions to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Shared Key (FWSM 3.1.1+ only) The options in this section let you encrypt the communications between the active and standby devices by providing a shared encryption key. ⚠ **Caution** All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If this device is used to terminate VPN tunnels, this information includes any user names, passwords and shared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communications with a shared key.	
Shared Key Confirm	Enter any string of characters up to 63 numbers, letters and punctuation characters. This string is used to generate the encryption key. Re-enter this string the **Confirm** field. If you select **HEX**, the entry in the Shared Key and Confirm fields must be exactly 32 hexadecimal characters (0-9, a-f).
Interface Configuration This table is presented on the Failover page for devices operating in single-context mode, or for individual security contexts only. The table lists all available named interfaces. To enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23. Select or deselect **Monitor this interface for failure**.	

Advanced Settings Dialog Box

The Advanced Settings dialog box lets you configure additional failover settings for the selected FWSM.

Note The following reference table describes all fields that can be presented in the Advanced Settings dialog box. The fields actually presented depend on operating mode (routed or transparent) and whether the device is hosting single or multiple contexts.

Navigation Path

You can access the Advanced Settings dialog box by clicking the Settings button on the Failover Page (FWSM), page 50-13.

Related Topics

- Failover Policies, page 50-10

Field Reference

Table 50-4 Advanced Settings Dialog Box

Element	Description
Interface Policy	
Select a failed-interfaces option and provide an appropriate value.	
Number of failed interfaces	When the number of failed monitored interfaces exceeds this value, the security appliance fails over. Valid values range from 1 to 250.
Percentage of failed interfaces	When the number of failed monitored interfaces exceeds this percentage, the security appliance fails over.
Failover Poll Time	
These fields define how often hello messages are sent on the failover link, and how long to wait before testing the peer for failure if no hello messages are received.	
Unit Failover	The amount of time between hello messages between failover units. Enter a value between 1 and 15 seconds, or if **msec** is checked, between 500 and 999 milliseconds.
Unit Hold Time	The amount of time to wait for a hello message on the failover link, after which the unit begins testing for peer failure. Enter a value between 3 and 45 seconds. This value must be at least three times the Unit Failover value.
Monitored Interface	The amount of time between polls among interfaces. Enter a value between 3 and 15 seconds.
MAC Address Mapping	
In Active/Standby mode, this table lists interface-virtual MAC address mappings. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, which are described in Using Tables, page 1-48. To add or edit interface mappings, click the Add Row or Edit Row button to open the Add/Edit Interface MAC Address Dialog Box, page 50-23.	

Table 50-4 Advanced Settings Dialog Box (continued)

Element	Description
Failover Groups	
	In Active/Active mode, this table lists both failover groups. To edit failover parameters for either group, select it in the list and click the Edit Row button to open the Edit Failover Group Dialog Box, page 50-25.
Bridge Group Configuration	
	In single-context transparent mode, this table lists all currently defined bridge groups (see Managing Device Interfaces, Hardware Ports, and Bridge Groups, page 46-26). To add a standby IP address to a bridge group, select it in the list and click the Edit Row button to open the Edit Failover Bridge Group Configuration Dialog Box, page 50-17.

Edit Failover Bridge Group Configuration Dialog Box

Use this dialog box to add a standby IP address to a failover bridge group.

Navigation Path

You can access the Edit Failover Bridge Group Configuration dialog box as follows:

- On the Failover page presented for an individual security context in transparent mode on an ASA.
- From the Bridge Group Configuration table in the Advanced Settings Dialog Box, page 50-16 presented by an FWSM in transparent mode.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13

Field Reference

Table 50-5 Edit Failover Bridge Group Configuration Dialog Box

Element	Description
Name	Identifies the bridge group; not editable.
IP Address	Identifies the IP address assigned to the bridge group; not editable.
Network Mask	Identifies the subnet mask for the IP Address; not editable.
Standby Address	Enter the IP address of the standby bridge group; this address must be on the same subnet as the primary address.

Failover Page (ASA/PIX 7.0+)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Use the Failover page to configure basic failover settings for ASA and PIX 7.0+ security devices

Note The features and options presented on the Failover page vary according to type of device selected, operating system version, firewall mode (routed or transparent), and security contexts (single or multiple). Thus, some of the elements described in the following table may not appear on the Failover page for your currently selected device.

Navigation Path

Select an ASA or PIX 7.0+ in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

Field Reference

Table 50-6 Failover Page (ASA/PIX 7.0+)

Element	Description
Failover Method	Choose the type of failover link: **Serial Cable** or **LAN Based**. If you choose Serial Cable, ensure the physical cable is connected to both devices. **Note** This option is available only on PIX devices.
Enable Failover	Check this box to enable failover on this device. Ensure that both devices have the same software version, activation key type, flash memory, and RAM. On PIX devices with LAN Based chosen as the Failover Method, and on all ASAs, you must next configure the logical LAN Failover interface and, optionally, the stateful failover interface.
Bootstrap button	Click to display the Bootstrap Configuration for LAN Failover dialog box. See Bootstrap Configuration for LAN Failover Dialog Box, page 50-26 for more information.
Settings button	Click to display the Settings Dialog Box, page 50-21, used to define when failover should occur.
Timeout	The failover Timeout specifies the amount of time after a system boots or becomes active that "nailed" sessions are accepted; used in conjunction with static translation rules (see Static Rules Tab, page 24-26 for more information). Enter a value in this field to specify the failover reconnect timeout value for asymmetrically routed sessions. The value is in the form hh:mm:ss (hours:minutes:seconds); both minutes and seconds are optional. Valid values for the number of hours are -1 to 1193; the default value is 0, which means connections cannot be re-established. Setting this value to -1 disables the timeout, allowing reconnections after any amount of time.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Configuration	
This section is presented only for devices operating in multiple-context mode.	
Active/Active	In an Active/Active failover configuration, both security appliances inspect network traffic, on a per-context basis. That is, for each context, one of the appliances is the active device, while the other is the standby device. To enable Active/Active failover on the security appliance, you must assign the security contexts to one of two failover groups. A failover group is a simply a logical group of one or more security contexts. You should specify failover group assignments on the unit that will have failover group 1 in the active state. The admin context is always a member of failover group 1. Any unassigned security contexts are also members of failover group 1 by default. See Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 for information about assigning a context to a failover group.
Active/Standby	In an Active/Standby configuration, the active security appliance handles all network traffic passing through the failover pair. The standby security appliance does not handle network traffic until a failure occurs on the active security appliance. Whenever the configuration of the active security appliance changes, it sends configuration information over the failover link to the standby security appliance. When a failover occurs, the standby security appliance becomes the active unit. It assumes the IP and MAC addresses of the previously active unit. Because the other devices on the network do not see any changes in the IP or MAC addresses, ARP entries do not change or time out.
LAN Failover	
Interface	Choose the interface to use as the failover link; all interfaces available on the device are listed. When configured for failover, the interface is directly connected to the standby device. **Note** You can choose an EtherChannel interface as the failover link. As with any other type of interface assigned as a failover link, the EtherChannel interface cannot be named, and none of the EtherChannel's member interfaces can be named. Further, while being used as an active failover link, changes to the interface configuration are not allowed. Refer to Configuring EtherChannels, page 46-9 for more information.
Logical Name	Enter a logical name for the failover interface.
Active IP Address	Specify the active IP address for this interface.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Standby IP Address	Specify a standby IP address for this interface. To receive packets from both units in a failover pair, standby IP addresses need to be configured on all interfaces. The Standby IP address is used on the security appliance that is currently the standby unit, and it must be in the same subnet as the active IP address.
Subnet Mask	Enter the Subnet Netmask for the active and standby IP addresses.
Stateful Failover (Optional) To configure Stateful Failover, page 50-4, provide the following parameters.	
Interface	Choose the interface to use for the stateful failover link; all interfaces available on the device are listed. **Note** You can choose an EtherChannel interface as the stateful failover link. As with any other type of interface assigned as a failover link, the EtherChannel interface cannot be named, and none of the EtherChannel's member interfaces can be named. Further, while being used as an active failover link, changes to the interface configuration are not allowed. Refer to Configuring EtherChannels, page 46-9 for more information.
Logical Name	Enter the logical name of the interface on the active firewall device to communicate with standby device for failover. When configured for stateful failover, the interface is directly connected to the standby device.
Active IP Address	Specify the IP address of the active interface.
Standby IP Address	Specify the IP address of the standby interface.
Subnet Mask	Enter the Subnet Netmask for the active and standby IP addresses.
Enable HTTP Replication	When selected, active HTTP sessions are copied to the standby firewall. Otherwise, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link.
Key The options in this section let you encrypt the communications between the active and standby devices. Select the type and provide a string of characters to produce the shared encryption key. ⚠ **Caution** All information sent over the failover and Stateful Failover links is sent in clear text unless you secure the communication with a failover key. If this device is used to terminate VPN tunnels, this information includes any user names, passwords and shared keys used for establishing the tunnels. Transmitting this sensitive data in clear text could pose a significant security risk. We recommend securing the failover communications with a shared key.	
Any string HEX	If you select **Any string**, the entry in the Shared Key field can be any combination of up to 63 numbers, letters and punctuation characters. This string is used to generate the encryption key. If you select **HEX**, the entry in the Shared Key and Confirm fields must be exactly 32 hexadecimal characters (0-9, a-f). This string is used as the encryption key.

Table 50-6 Failover Page (ASA/PIX 7.0+) (continued)

Element	Description
Shared Key Confirm	Enter any string of characters appropriate to the selected key type: Any string or HEX. Re-enter the string the **Confirm** field.

Interface Configuration
(in some instances, labeled Monitor Interface Configuration)

This table is presented on the Failover page for ASA 8.4.1+ devices operating in single-context, transparent mode, and for individual contexts on PIX/ASA devices. Otherwise, it appears in the Settings Dialog Box, page 50-21.

The table lists all available named interfaces. To enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23. Select or deselect **Monitor this interface for failure**.

Settings Dialog Box

The Settings dialog box lets you define criteria for when failover should occur on the selected ASA or PIX 7.x appliance.

Navigation Path

You can access the Settings dialog box by clicking the Settings button on the Failover Page (ASA/PIX 7.0+), page 50-17.

Note The following reference table presents all fields that can be presented in the Settings dialog box. The fields actually presented depend on operating mode (routed or transparent) and whether the device is hosting single or multiple contexts.

Related Topics

- Failover Policies, page 50-10
- Edit Failover Interface Configuration Dialog Box, page 50-23
- Add/Edit Interface MAC Address Dialog Box, page 50-23
- Bootstrap Configuration for LAN Failover Dialog Box, page 50-26

Field Reference

Table 50-7 Settings Dialog Box

Element	Description
Interface Policy	
Number of failed interfaces	When the number of failed monitored interfaces exceeds this value, the security appliance fails over. The range is between 1 and 250 failures.
Percentage of failed interfaces	When the number of failed monitored interfaces exceeds this percentage, the security appliance fails over.

Table 50-7 Settings Dialog Box (continued)

Element	Description
Failover Poll Time	
Unit Failover	The amount of time between hello messages among units. The range is between 1 and 15 seconds, or between 200 and 999 milliseconds if the **Change units to msec** option is checked.
Unit Hold Time	Sets the time during which a unit must receive a hello message on the failover link, or the unit begins the testing process for peer failure. The range is between 3 and 45 seconds, or between 800 and 999 milliseconds if the **msec** option is checked. You cannot enter a value that is less than three times the Unit Failover value.
Monitored Interface	The amount of time between polls among interfaces. The range is between 3 and 15 seconds, or between 500 and 999 milliseconds if the msec option is checked.
Interface Hold Time	Sets the time during which a data interface must receive a hello message, after which the peer is declared failed. Valid values are from 5 to 75 seconds. This value must be at least five times the Unit Failover value.
Link State Interval	Sets the interval after which each ASA in a failover pair checks the link state of its interfaces. By default the link state interval value is 500msec. You can customize the polltime; for example, if you set the polltime to 300 msec, the ASA can detect an interface failure and trigger failover faster. Valid range is between 300 and 799 milliseconds. **Note** The Link State Interval is available for ASA 9.7.1 and higher.
Failover Groups In Active/Active mode, this table lists both failover groups. To edit failover parameters for either group, select it in the list and click the Edit Row button to open the Edit Failover Group Dialog Box, page 50-25.	
MAC Address Mapping In Active/Standby mode, this table lists interface-virtual MAC address mappings. This is a standard Security Manager table, with Add Row, Edit Row and Delete Row buttons, which are described in Using Tables, page 1-48. To add or edit interface mappings, click the Add Row or Edit Row button to open the Add/Edit Interface MAC Address Dialog Box, page 50-23.	
Monitor Interface Configuration In single-context mode, this table lists all available named interfaces. To define a Standby IP address for, and enable or disable monitoring of an interface, select it in the list and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23.	
Management IP Address In single-context transparent mode, this section presents the management IP address and netmask defined for the device (on the Management IP Page, page 47-10); you cannot change these values.	
Standby	Enter the management IP address of the standby unit; this address must be on the same subnet as the primary address.

Add/Edit Interface MAC Address Dialog Box

The Add/Edit Interface MAC Address dialog box lets you define virtual MAC addresses for a physical interface on ASA, FWSM 3.x and PIX 7.x security appliances that are configured for failover (not available on ASA 5505 devices).

In Active/Standby failover, the MAC addresses for the primary unit are always associated with the active IP addresses. If the secondary unit boots first and becomes active, it uses the burned-in MAC address for its interfaces. When the primary unit comes online, the secondary unit obtains the MAC addresses from the primary unit. This change can disrupt network traffic. You can configure virtual MAC addresses for each interface to ensure that the secondary unit uses the correct MAC addresses when it is the active unit, even if it comes online before the primary unit. If you do not specify virtual MAC addresses, the failover pair uses the burned-in MAC addresses.

Note You cannot configure a virtual MAC address for the failover or Stateful Failover links. The MAC and IP addresses for those links do not change during failover.

Navigation Path

You can open the Add/Edit Interface MAC Address dialog box from the Settings Dialog Box, page 50-21.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Edit Failover Group Dialog Box, page 50-25

Field Reference

Table 50-8 Add/Edit Interface MAC Address Dialog Box

Element	Description
Physical Interface	Choose the physical interface on which failover virtual MAC addresses are to be configured.
MAC Address	
Active Interface	Enter a virtual MAC address for the active interface in hexadecimal format (for example, 0023.4567.89ab).
Standby Interface	Enter a virtual MAC address for the standby interface in hexadecimal format (for example, 0023.4567.89ab).

Edit Failover Interface Configuration Dialog Box

Use the Edit Failover Interface Configuration dialog box to define a standby IP address for an interface, and to specify whether the status of the interface should be monitored.

Note A failover interface cannot be configured for PPPoE.

Navigation Path

You can access the Edit Failover Interface Configuration dialog box from the Settings Dialog Box, page 50-21 (ASA/PIX 7.0+), Advanced Settings Dialog Box, page 50-16 (FWSM), and from the Failover page itself for ASA 8.4.1+ devices operating in single-context transparent mode, and for individual ASA/PIX security contexts.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13
- Edit Failover Group Dialog Box, page 50-25

Field Reference

Table 50-9 Edit Failover Interface Configuration Dialog Box

Element	Description
Interface Name	The name of the interface; read-only.
Active IP Address	The active IP address of this interface; read-only. This field is blank if an IP address has not been assigned to the interface; for example if DHCP is enabled on the interface.
Mask	The subnet mask for the active IP address; read-only. This field is blank if an IP address has not been assigned to the interface; for example, if DHCP is enabled on the interface.
Standby IP Address	Specify the IP address of the corresponding interface on the standby failover unit. This field does not appear if an IP address has not been assigned to the interface.
Monitor this interface for failure	Specifies whether this interface is monitored for failure: check this box to enable monitoring. The number of interfaces that can be monitored for the security appliance is 250. Hello messages are exchanged between the security appliance failover pair during every interface poll time period. The failover interface poll time is 3 to 15 seconds. For example, if the poll time is set to 5 seconds, testing begins on an interface if 5 consecutive hellos are not heard on that interface (25 seconds). Monitored failover interfaces can have the following status: • Unknown—Initial status. This status can also mean the status cannot be determined. • Normal—The interface is receiving traffic. • Testing—Hello messages are not heard on the interface for five poll times. • Link Down—The interface is administratively down. • No Link—The physical link for the interface is down. • Failed—No traffic is received on the interface, yet traffic is heard on the peer interface.

Table 50-9 Edit Failover Interface Configuration Dialog Box (continued)

Element	Description
ASR Group Number	If this interface is part of an asymmetric routing group, provide its ASR group number. Valid values for ASR group numbers are 1 through 32. Stateful failover must be enabled for asymmetric routing support to function properly between units in failover configurations.

Edit Failover Group Dialog Box

Use the Edit Failover Group dialog box to configure failover parameters for groups of security contexts in an Active/Active failover configuration. See Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9, or Add/Edit Security Context Dialog Box (FWSM), page 59-8, for information about assigning a context to a failover group.

Navigation Path

You can access the Add Failover Group dialog box from the PIX/ASA Settings Dialog Box, page 50-21, or the FWSM Advanced Settings Dialog Box, page 50-16.

Related Topics

- Failover Policies, page 50-10
- Failover Page (ASA/PIX 7.0+), page 50-17
- Failover Page (FWSM), page 50-13

Field Reference

Table 50-10 Edit Failover Group Dialog Box

Element	Description
Preferred Role	Specifies the unit in the failover pair, primary or secondary, on which this failover group appears in the active state when both units start up simultaneously, or when the Preempt option is selected. Choose **Primary** or **Secondary**. You can have both failover groups in the active state on a single unit in the pair; however, a more typical configuration is to assign each failover group a different role to make each one active on a different unit, balancing the traffic across the devices.
Poll time interval for monitored interfaces	Specify the amount of time between polling of monitored interfaces. Valid values are from 3 to 15 seconds (or 500 to 999 milliseconds if **msec** is checked).
Hold Time	Specify the time period within which the group must receive a hello message, after which the other group is declared failed. Valid values are from 5 to 75 seconds.
Preempt after Reboot	Specifies the number of seconds that the preferred failover device should wait after rebooting before taking over as the active unit for this failover group. Valid values are from 0 to 1200 seconds.

Table 50-10 Edit Failover Group Dialog Box (continued)

Element	Description
Enable HTTP Replication	Indicates whether active HTTP sessions are copied to the standby device for this failover group as part of Stateful failover. If you do not allow HTTP replication, HTTP connections are disconnected at failover. Disabling HTTP replication reduces the amount of traffic on the state link. This setting overrides the HTTP replication setting on the Failover page.
Failover Criteria	Select a failed-interfaces criterion for this group and specify the appropriate value: • **Number of failed interfaces** – When this number of interfaces have failed, failover is triggered. Valid values are 1 to 250. • **Percentage of failed interfaces** – When this percentage of the total number of interfaces have failed, failover is triggered. Valid values are 1 to 100.
MAC Address Mapping This table displays interfaces to which active and standby MAC addresses are mapped.	

Failover Page (Security Context)

The Failover page for individual ASA and PIX 7.0+ security contexts presents the **Interface Configuration** table, which lists all available named interfaces.

You can select an interface in the table and click the Edit Row button to open the Edit Failover Interface Configuration Dialog Box, page 50-23, where you can specify a standby IP address and an ASR group number, and enable or disable monitoring of the interface.

For individual transparent-mode contexts on ASA 8.4.1+ devices, the Failover page also presents the **Bridge Group Configuration** table, which lists all currently defined failover bridge groups.

You can select an entry in the table and click the Edit Row button to open the Edit Failover Bridge Group Configuration Dialog Box, page 50-17, where you can specify a standby IP address for the selected bridge group.

Navigation Path

Select a security context in Device View and then select **Platform > Device Admin > Failover** from the Device Policy selector.

Related Topics

- Understanding Failover, page 50-1
- Failover Policies, page 50-10
- About Bridging on Firewall Devices, page 47-1

Bootstrap Configuration for LAN Failover Dialog Box

The Bootstrap Configuration for LAN Failover dialog box provides you with bootstrap configuration that can be applied to the primary and secondary devices in a LAN failover configuration.

Navigation Path

You can access the Bootstrap Configuration for LAN Failover dialog box from the Failover page. For more information about the Failover page, see:

- Failover Page (PIX 6.3), page 50-10
- Failover Page (FWSM), page 50-13
- Failover Page (ASA/PIX 7.0+), page 50-17

Related Topics

- Failover Policies, page 50-10
- Additional Steps for an Active/Standby Failover Configuration, page 50-9

Field Reference

Table 50-11 Bootstrap Configuration for LAN Failover Dialog Box

Element	Description
Primary	Contains the bootstrap configuration for the primary device. Open a console connection to the primary device and then paste this configuration to activate failover on the device.
Secondary	Contains the bootstrap configuration for the secondary device. After the primary device becomes active, open a console connection to the secondary device and then paste this configuration to activate failover on the device.

Note For Active/Active Failover, the bootstrap configurations are only applied to the system contexts of the respective failover peer devices.

Configuring Hostname, Resources, User Accounts, and SLAs

The following topics describe configuring the host name on a security appliance, defining and managing Resource classes on Firewall Services Modules (FWSMs) in multiple-context mode, managing user accounts in the Local user database, and monitoring service level agreements (SLAs) to perform route tracking.

This chapter contains the following topics:

- Hostname Page, page 51-1
- Resource Management on Multi-context FWSMs, page 51-2
- Configuring User Accounts, page 51-7
- Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8

Hostname Page

You can use the Hostname page to specify a host name for your security device, and to specify a default domain. After the configuration file is deployed, the device uses this domain name when you do not enter a fully-qualified domain name in other commands. It also uses this domain name in RSA key generation.

The device appends this domain name to unqualified names. For example, if you set the domain name to “example.com,” and specify a syslog server by the unqualified name “jupiter,” the security appliance completes the name to “jupiter.example.com.”

When you set a host name for the security appliance, that name appears in the command line prompt. If you establish sessions to multiple devices, the host name helps you keep track of where you enter commands. The default host name depends on your platform.

In multiple-context mode, you can specify a domain name for each context, as well as the system execution space. The host name you specify in the system execution space appears in the command line prompt for all contexts. The host name that you optionally set within a context does not appear in the command line, but can be used by the banner command `$(hostname)` token.

Navigation Path

In Device View, select a security device and then select **Platform > Device Admin > Hostname** from the Device Policy selector.

Field Reference

Table 51-1 Hostname Page

Element	Description
Host Name	Enter a unique device name to help you differentiate among devices; for example, *PIX-510-A*. **Note** We recommend that you use a unique host name for each device you manage. The device name can be up to 63 alphanumeric (U.S. English) characters and can include any of the following special characters: ` ( ) + - , . / : =.
Domain Name	Optionally, enter a valid Domain Name System (DNS) domain name for the device; for example, *cisco.com*.

Resource Management on Multi-context FWSMs

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

By default, all security contexts on a multiple-context Firewall Services Module (FWSM) have unlimited access to the resources of the FWSM, except where maximum limits per context are enforced. However, if you find that one or more contexts use too many resources, and they cause other contexts to be denied connections, for example, then you can configure resource management to limit the use of resources per context.

Note The FWSM does not limit the bandwidth per context; however, the switch containing the FWSM can limit bandwidth per VLAN. See the switch documentation for more information.

The FWSM manages resources by assigning contexts to resource classes. Each context uses the resource limits set by its class. When you create a class, the FWSM does not set aside a portion of the resources for each context assigned to the class; rather, the FWSM sets the maximum limit for a context. If you oversubscribe resources, or allow some resources to be unlimited, a few contexts can "use up" those resources, potentially affecting service to other contexts.

You can set the limit for all resources together as a percentage of the total available for the device. Also, you can set the limit for individual resources as a percentage or as an absolute value.

You can oversubscribe the FWSM by assigning more than 100 percent of the resources across all contexts. For example, you can set up a class to limit connections to 20 percent per context, and then assign 10 contexts to the class for a total of 200 percent. If contexts concurrently use more than the system limit, then each context gets less than the 20 percent you intended.

The FWSM also lets you assign unlimited access to one or more resources in a class, instead of a percentage or absolute number. When a resource is unlimited, contexts can use as much of the resource as the system has available. For example, contexts A, B, and C are assigned to class "Onepercent," which limits each class member to one percent of the system inspections per second, for a total of three percent; but the three contexts are currently only using two percent combined. On the other hand, class "Nolimit" has unlimited access to inspections. The contexts in Nolimit can use more than the 97 percent of "unassigned" inspections; they can also use the one percent of inspections not currently in use by

contexts A, B, and C, even if that means that contexts A, B, and C are unable to reach their three percent combined limit. Setting unlimited access is similar to oversubscribing the FWSM, except that you have less control over how much you oversubscribe the system.

Default Class

All contexts belong to the default class if they are not assigned to another class; you do not have to actively assign a context to the default class.

If a context belongs to a class other than the default class, those class settings always override the default class settings. However, if the other class has any settings that are not defined, then the member context uses the default class for those limits. For example, if you create a class with a two percent limit for all concurrent connections, but no other limits, then all other limits are inherited from the default class. Conversely, if you create a class with a two percent limit for all resources, the class uses no settings from the default class.

As initially configured, the default class provides unlimited access to resources for all contexts, except for the following limits, which are by default set to the maximum allowed per context:

- Telnet sessions – 5 sessions
- SSH sessions – 5 sessions
- IPSec sessions – 5 sessions
- MAC addresses – 65,535 entries

Note that you can edit the default class.

Related Topics

- Resources Page, page 51-3
- Add/Edit Security Context Dialog Box (FWSM), page 59-8

Resources Page

Use the Resources page to configure and manage resource-management classes.

The table on this page lists all currently defined resource classes. Use the buttons below the table to manage this list:

- Add Row – Opens the Add Resource dialog box, where you can define a new class, and assign it to security contexts. See Add and Edit Resource Dialog Boxes, page 51-4 for more information.
- Edit Row – For the currently selected row, opens the Edit Resource dialog box, so you can edit that class and its context assignments. See Add and Edit Resource Dialog Boxes, page 51-4 for more information.
- Delete Row – Deletes the currently selected row(s); confirmation may be required.

Navigation Path

In Device View, select the system context of an ASA or FWSM in multiple-context mode, and then select **Platform > Device Admin > Resources** from the Device Policy selector.

Related Topics

- Resource Management on Multi-context FWSMs, page 51-2

Add and Edit Resource Dialog Boxes

Use the Add Resource and Edit Resource dialog boxes to add or edit resource classes and assignments for FWSM and ASA security contexts.

Except for their titles, both dialog boxes are identical; the following descriptions apply to both.

Navigation Path

You can access the Add Resource and Edit Resource dialog boxes from the Resources Page, page 51-3.

Related Topics

- Resource Management on Multi-context FWSMs, page 51-2

Field Reference

Table 51-2 Add and Edit Resource Dialog Boxes

Element	Description
Class Name	Enter a name for this class; can be a string of up to 20 alphanumeric characters, and may include any of the following special characters: ` ( ) + - , . / : =.
Limits Tab	
Note For the following Limits, if you do not specify a value for a particular limit, the limit is inherited from the default class. If the default class does not set that limit, the limit inherits the system limit. Also, any value you enter is considered to be that rate *per second*, unless you also check the related **percent** box, in which case the value is that percentage of the total resource.	
TCP or UDP Connections	Sets a Rate Limit for TCP or UDP connections between any two hosts, including connections between one host and multiple other hosts. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 102400, or you can assign more than 100 percent if you want to oversubscribe the device.
Inspections (Fixups)	Sets a Rate Limit for application inspections. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 10000 per second, or you can assign more than 100 percent if you want to oversubscribe the device.
Syslog Messages	Sets a Rate Limit for system log messages. You can set the limit as an absolute value, or you can assign more than 100 percent if you want to oversubscribe the device. The FWSM can support 30,000 messages per second for messages sent to the FWSM terminal or buffer. If you send messages to a syslog server, the FWSM supports 25,000 per second.

Table 51-2 Add and Edit Resource Dialog Boxes (continued)

Element	Description
Connections	Sets the Absolute Limit for concurrent TCP or UDP connections. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 999900, or you can assign more than 100 percent if you want to oversubscribe the device. **Note** For concurrent connections, the FWSM allocates half of the limit to each of two network processors (NPs) that accept connections. Typically, the connections are divided evenly between the NPs. However, in some circumstances, the connections are not evenly divided, and the maximum connection limit could be reached on one NP before reaching the maximum on the other. In this case, the maximum connections allowed is less than the limit you set. The NP distribution is controlled by the switch, based on a distribution algorithm. You can adjust this algorithm on the switch (see Customizing the FWSM Internal Interface on cisco.com), or you can adjust the connection limit upward to account for the inequity.
Hosts	Sets the limit for concurrent hosts that can connect through the FWSM. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 262144, or you can assign more than 100 percent if you want to oversubscribe the device.
IPsec Sessions	Sets the limit for IPsec sessions. You can set the limit as an absolute value by entering an integer between 1 and 5, or you can assign more than 100 percent if you want to oversubscribe the device. The system allows a maximum of 10 concurrent sessions divided between all contexts.
SSH Sessions	Sets the limit for SSH sessions. You can set the limit as an absolute value by entering an integer between 1 and 5, or you can assign more than 100 percent if you want to oversubscribe the device. The system allows a maximum of 100 concurrent sessions divided between all contexts.
Telnet Sessions	Sets the limit for concurrent Telnet sessions. You can set the limit as an absolute value by entering an integer between 1 and 5, or you can assign more than 100 percent if you want to oversubscribe the device. The system allows a maximum of 100 concurrent sessions divided between all contexts.
NAT Translations	Sets the limit for concurrent address translations. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 266144, or you can assign more than 100 percent if you want to oversubscribe the device.
MAC Address	(Transparent mode only) Sets the limit for concurrent MAC address entries allowed in the MAC address table. You can set the limit as an absolute value by entering an integer between 0 (system limit) and 65535, or you can assign more than 100 percent if you want to oversubscribe the device.

Table 51-2 Add and Edit Resource Dialog Boxes (continued)

Element	Description
ASDM	Sets the limit for ASDM management sessions (the default is 5). You can set the limit as an absolute value by entering an integer between 1 and 5, or you can enter a percentage between 3.0 and 15.0. The system allows a maximum of 80 concurrent sessions divided between all contexts. ASDM sessions use two HTTPS connections: one for monitoring that is always present, and one for making configuration changes that is present only when you make changes. For example, the system limit of 80 ASDM sessions represents a limit of 160 HTTPS sessions, divided between all contexts.
Other VPN	Sets the limit for Site-to-site VPN sessions. You cannot oversubscribe this resource; all context assignments combined cannot exceed the model limit. The sessions you assign for this resource are guaranteed to the context.
Other VPN Burst	Sets the limit for the number of site-to-site VPN sessions allowed beyond the amount assigned to a context with vpn other. For example, if your model supports 5000 sessions, and you assign 4000 sessions across all contexts with vpn other, then the remaining 1000 sessions are available for other vpn burst. Unlike other vpn, which guarantees the sessions to the context, other vpn burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
Note The maximum value for Anyconnect VPN and Anyconnect VPN Burst depends on ASA licenses. Cisco Security Manager cannot validate the values entered for Anyconnect VPN and Anyconnect VPN Burst. Therefore, the user should make sure that the values for Anyconnect VPN and Anyconnect VPN Burst are within the maximum values; else it results in a deployment error. To find the maximum value, telnet into ASA and execute the **show version** command. The Total VPN Peers value corresponds to the maximum value.	
Anyconnect VPN	AnyConnect peers. You cannot oversubscribe this resource; all context assignments combined cannot exceed the model limit. The peers you assign for this resource are guaranteed to the context.
Anyconnect VPN Burst	The number of AnyConnect sessions allowed beyond the amount assigned to a context with AnyConnect. For example, if your model supports 5000 peers, and you assign 4000 peers across all contexts with AnyConnect, then the remaining 1000 sessions are available for AnyConnect Burst. Unlike AnyConnect, which guarantees the sessions to the context, AnyConnect Burst can be oversubscribed; the burst pool is available to all contexts on a first-come, first-served basis.
Storage	Beginning with version 4.12, Security Manager enables you to enter the storage size or select Default. This feature is available for ASA version 9.6(2) or later. The limit is set in MB. The default limit is 100% of the disk configured since this storage cannot span multiple disks.
All Resources Limit	Sets a limit for all resources. If you also set the limit for a specific resource, then that limit overrides the limit you set here for all resources. You can set the limit as a percentage, or as unlimited by setting the value to **0** (when **percent** is not checked). You cannot set any other absolute value. You can assign more than 100 percent if you want to oversubscribe the device.

Table 51-2 Add and Edit Resource Dialog Boxes (continued)

Element	Description
Contexts Tab	
Available Contexts	Lists all contexts available for class assignments; contexts which already have class assignments are not displayed. Select one or more contexts and click the >> button to add the contexts to the Selected Contexts list.
Selected Contexts	Lists all contexts assigned to this class. Select one or more contexts and click the << button to return the contexts to the Available Contexts list.

Configuring User Accounts

The User Accounts page lets you manage the Local user database. User accounts in the Local database can be used in conjunction with the Authentication, Authorization, Accounting (AAA) functions to determine "who is allowed to do what" on a device. Refer to About AAA on Security Devices, page 48-1 for more information.

The table on this page lists all currently defined Local user accounts, showing for each, the name of the user and the assigned privilege level. For a detailed explanation of these fields, see Add/Edit User Account Dialog Boxes, page 51-7.

- To add a user account, click the Add Row button.
- To edit the settings for an account, select it and click the Edit Row button.
- To delete a user account, select it and click the Delete Row button.

Navigation Path

- (Device view) Select **Platform > Device Admin > User Accounts** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > User Accounts** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Local Database, page 48-3
- Preparing for AAA, page 48-2

Add/Edit User Account Dialog Boxes

Use the Add and Edit User Account dialog boxes to add a local user account or to modify an existing user account.

Navigation Path

You can access the Add and Edit User Account dialog boxes from the User Accounts page, as described in Configuring User Accounts, page 51-7.

Field Reference

Table 51-3 Add/Edit User Account Dialog Boxes

Element	Description
Username	Enter a name for this user account: must be at least four characters; the maximum is 64 characters. Entries are case-sensitive.
Password	
Password as encrypted	Select Plain Text or Encrypted.
Password encrypt type	Select MD5 or PBKDF2.
Password	Enter a unique password for this user account. Entries are case-sensitive. **Note** To protect security, we recommend a password length of at least 8 characters. **Note** For Plain Text passwords: – The length of MD5 password should be three to 32 characters. – The length of PBKDF2 password should be 33 to 127 characters.
Confirm	Re-enter the user password to confirm it.
Privilege Level	Choose a privilege level for this user; defines local command authorization. The range is 0 (lowest) to 15 (highest). The default privilege level is 2.

Monitoring Service Level Agreements (SLAs) To Maintain Connectivity

You can configure ASA or PIX devices that run version 7.2 or higher to perform route tracking by monitoring service level agreements. By monitoring the connectivity to a device on another network, you can track the availability of a primary route and install a backup route if the primary route fails. For example, you can define a default route to an Internet service provider (ISP) gateway and a backup default route to a secondary ISP in case the primary ISP becomes unavailable. This technique, called Dual ISP, provides security appliances with a form of high availability, which is a vital part of providing customers with the services to which they are entitled.

Without route tracking, there is no inherent mechanism for determining if the route is up or down. A static route remains in the routing table even if the next hop gateway becomes unavailable, and is removed only if the associated interface on the security appliance goes down.

The security appliance performs route tracking by associating a route with a monitoring target that you define in an SLA monitor policy object. It monitors the target using ICMP echo requests, according to the parameters configured in the object. If an echo reply is not received within a specified time period, the SLA monitor is considered down and the associated route is removed from the routing table. A previously configured backup route is used in place of the removed route.

SLA monitoring jobs start immediately after deployment and continue to run unless you remove the SLA monitor from the device configuration (that is, they do not age out).

Related Topics

- Configuring Static Routes, page 56-128
- Configuring Firewall Device Interfaces, page 46-3
- Creating Policy Objects, page 6-9

This section contains the following topics:

- Creating Service Level Agreements, page 51-9

Creating Service Level Agreements

The following procedure explains how to configure SLA monitor objects and associate them with routes and interfaces in an ASA or PIX configuration.

Related Topics

- Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8
- Configuring Static Routes, page 56-128
- Configuring Firewall Device Interfaces, page 46-3
- Creating Policy Objects, page 6-9

Step 1 Create the SLA monitor policy object:

a. Select **Manage > Policy Objects** to open the Policy Object Manager (see Policy Object Manager, page 6-4) and select **SLA Monitors** from the table of contents.

Tip You can also create SLA monitor objects when defining policies that use this object type. For more information, see Selecting Objects for Policies, page 6-2.

b. Right-click in the work area and select **New Object** to open the Add SLA Monitor dialog box. For more information, see Configuring SLA Monitor Objects, page 51-10.

c. The monitoring options are appropriate for most connections, so you need only configure the following:

- Name—The name of the object.
- SLA Monitor ID—An identifying number for the monitoring process. The number must be unique within a device configuration.
- Monitored Address—The address that you are monitoring. When you select a monitoring target, make sure that it can respond to ICMP echo requests (pings). The target can be any network address that you choose, but consider the use of:
- The ISP gateway address.
- The next hop gateway address (if you are concerned about the availability of the ISP gateway).
- A server on the target network, such as an AAA server, with which the security appliance needs to communicate.
- A persistent network device on the destination network. (A desktop or notebook computer that gets shut down at night is not a good choice.)

- Interface—The interface name, or interface role that identifies an interface, that will be the source of the ICMP messages. The device pings the monitored address from this interface's IP address.

d. Click **OK** to save the object.

Step 2 Configure ASA/PIX policies to use the object to monitor routes. You can configure the following policies to monitor SLAs:

- **Platform > Routing > Static Route**—When you define a static route, you can select an SLA monitor object to do route tracking for the route. For more information, see Configuring Static Routes, page 56-128 and Add/Edit Static Route Dialog Box, page 56-130.
- **Interfaces**—When you define an interface that uses DHCP or PPPoE, you can configure the DHCP or PPPoE learned default routes to be tracked. For more information, see Device Interface: IP Type (PIX/ASA 7.0+), page 46-58.

Configuring SLA Monitor Objects

Use the Add or Edit SLA (Service Level Agreement) Monitor dialog box to create, edit, and copy SLA monitor objects. Each SLA monitor defines a connectivity policy to a monitored address and tracks the availability of a route to the address. The route is periodically checked for availability by sending ICMP echo requests and waiting for the response. If the requests time out, the route is removed from the routing table and replaced with a backup route.

You can configure SLA monitors only for security appliances running PIX/ASA version 7.2 or higher. SLA monitoring jobs start immediately after deployment and continue to run unless you remove the SLA monitor from the device configuration (that is, they do not age out).

For more information about configuring and using SLA monitor objects, see Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8.

Navigation Path

Select **Manage > Policy Objects**, then select **SLA Monitors** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8
- Policy Object Manager, page 6-4

Field Reference

Table 51-4 SLA Monitor Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.

Table 51-4 SLA Monitor Dialog Box (continued)

Element	Description
SLA Monitor ID	The ID number of the SLA operation. Values range from 1 to 2147483647. You can create a maximum of 2000 SLA operations on a device. Each ID number must be unique to the policy and the device configuration.
Monitored Address	The IP address that is being monitored for availability by the SLA operation. For recommendations on selecting an address to monitor, see Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8.
Interface	The source interface for all ICMP echo requests sent to the monitored address to test its availability. Enter the name of an interface or interface role, or click **Select** to select an it from a list or to create a new interface role.
Frequency	The frequency of ICMP echo request transmissions, in seconds. Values range from 1 to 604800 seconds (7 days). The default is 60 seconds. **Note** The frequency cannot be less than the timeout value; you must convert frequency to milliseconds to compare the values.
Threshold	The amount of time that must pass after an ICMP echo request before a rising threshold is declared, in milliseconds. Values range from 0 to 2147483647 milliseconds. The default is 5000 milliseconds. The threshold value is used only to indicate events that exceed the defined value. You can use these events to evaluate the proper timeout value. It is not a direct indicator of the reachability of the monitored address. **Note** The threshold value should not exceed the timeout value.
Time out	The amount of time that the SLA operation waits for a response to the ICMP echo requests, in milliseconds. Values range from 0 to 604800000 milliseconds (7 days). The default is 5000 milliseconds. If a response is not received from the monitored address within the amount of time defined in this field, the static route is removed from the routing table and replaced by the backup route. **Note** The timeout value cannot exceed the frequency value (adjust the frequency value to milliseconds to compare the numbers).
Request Data Size	The size of the ICMP request packet payload, in bytes. Values range from 0 to 16384 bytes. The default is 28 bytes, which creates a total ICMP packet of 64 bytes. Do not set this value higher than the maximum allowed by the protocol or the Path Maximum Transmission Unit (PMTU). For purposes of reachability, you might need to increase the default data size to detect PMTU changes between the source and the target. A low PMTU can affect session performance and, if detected, might indicate that the secondary path should be used.

Table 51-4 SLA Monitor Dialog Box (continued)

Element	Description
ToS	The type of service (ToS) defined in the IP header of the ICMP request packet. Values range from 0 to 255. The default is 0. This field contains information such as delay, precedence, reliability, and so on. It can be used by other devices on the network for policy routing and features such as committed access rate.
Number of Packets	The number of packets that are sent. Values range from 1 to 100. The default is 1 packet. **Tip** Increase the default number of packets if you are concerned that packet loss might falsely cause the security appliance to believe that the monitored address cannot be reached.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

CHAPTER 52

Configuring Server Access Settings on Firewall Devices

The Server Access section contains pages for configuring server access on firewall devices; Server Access is under Device Admin in the Device or Policy selector.

This chapter contains the following topics:

- AUS Page, page 52-1
- DHCP Relay Page, page 52-5
- DHCP Relay IPv6 Page, page 52-7
- Configuring DHCP Servers, page 52-10
- DNS Page, page 52-14
- Configuring DDNS, page 52-18
- NTP Page, page 52-21
- SMTP Server Page, page 52-22
- TFTP Server Page, page 52-23

AUS Page

The AUS page lets you configure remote updating of a security appliance from a server that supports the Auto Update specification. Auto Update applies configuration changes and software updates to the appliance automatically from the remote server.

Note The server you identify on this page must be the same server you identify in the Auto Update section of the Device Properties (from the Tools menu, choose Device Properties). The Device Properties information identifies the AUS server to which Security Manager sends configuration updates, whereas the information on this page defines for the server the device will contact for updates. Also, the Device Identity you provide in the Device Properties must match the Device ID on this page.

If you change AUS servers, note that the device will continue to use the AUS server defined in its current configuration until it receives a new configuration. Thus, you should change the AUS policy but deploy the configuration using the previous AUS server. After deployment is successful, change the Device Properties to point to the new server. For more information on deploying to AUS, see Deploying Configurations Using an Auto Update Server or CNS Configuration Engine, page 8-41.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > AUS** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > AUS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add and Edit Auto Update Server Dialog Boxes, page 52-3

Field Reference

Table 52-1 AUS Page

Element	Description
Auto Update Servers table	This table lists currently configured Auto Update servers. Use the buttons below the table to manage these entries. The entries are listed in order of precedence for contacting AUS servers. Use the Up and Down arrow buttons to change the ordering of the list by moving the selected entry up or down. Use the Add Row, Edit Row, and Delete Row buttons to add, edit or delete entries. Add Row opens the Add Auto Update Server dialog box, while Edit Row opens the Edit Auto Update Server dialog box for the selected row. See Add and Edit Auto Update Server Dialog Boxes, page 52-3 for information about these dialog boxes. **Note** The URL for contacting this AUS server is produced by concatenating the `Protocol`*://Username:Password@IP IP Address*(*:Port*)/*Path* provided in the Add/Edit Auto Update Server dialog boxes. The Port is included only if you entered a port number other than the default 443.
Device ID Type	Choose the method used for identifying this device to the AUS server: • **Host Name** – The host name of this device, as provided in the Device Properties window (Tools > Device Properties). • **Serial Number** – The serial number of this device. • **IP Address** – The IP address of the specified interface. When you choose this option, an Interface field appears; enter or Select the desired device interface. • **MAC Address** – The MAC address of the specified interface. When you choose this option, an Interface field appears; enter or Select the desired device interface. • **User Defined** – A unique user-specified ID is used. When you choose this option, a User Defined field appears; enter any alphanumeric string. Note that this string must also appear in the Device Identity field in the Device Properties window (Tools > Device Properties).

Table 52-1 AUS Page (continued)

Element	Description
Poll Type	Choose the method defining how often the AUS server is polled for updates: • **At Specified Frequency** – If you choose this option, the Poll Period field is displayed: – **Poll Period** – Specify the number of minutes the device waits between polls of the AUS server; valid values are 1 to 35791. • **At Scheduled Time** – If you choose this option, the following fields are displayed (available only on ASA/PIX devices running version 7.2 or later): – **Days of the week** – Select one or more days on which the device is to poll the AUS server. – **Polling Start Time in Hours** – The hour at which polling is to begin on the selected days; based on a 24-hour clock. – **Polling Start Time in Mins** – The minute within the chosen hour when polling is to begin. – **Enable Randomization of the Start Time** – Select this option to specify a random polling window; the Randomization Window field is enabled. **Randomization Window** – The maximum number of minutes the device can use to randomize the specified polling time; valid values are 1 to 1439.
Retry Count	The number of times the device will try to poll the AUS server for new information. Optional; if you enter zero or leave this field blank, the device will not retry after a failed poll attempt.
Retry Period	If Retry Count is not zero or blank, the number of minutes the device will wait to re-poll the AUS server if the previous attempt failed; valid values are 1 to 35791. If Retry Count is not zero or blank and you leave this field blank, the value defaults to five minutes.
Disable Device After:	Selecting this option ensures that if no response is received from the AUS server within the specified Timeout period, the security appliance will stop passing traffic. • **Timeout** – The number of minutes the firewall device will wait to timeout if no response is received from the AUS server.

Add and Edit Auto Update Server Dialog Boxes

Use the Add Auto Update Server dialog box to configure a new AUS server definition. The security appliance will automatically poll this server for image and configuration updates.

The Auto Update specification allows the Auto Update server to either push configuration information and send requests for information to the security appliance, or to pull configuration information by causing the security appliance to periodically poll the Auto Update server. The Auto Update server can

also send a command to the security appliance to send an immediate polling request at any time. Communication between the Auto Update server and the security appliance requires a communications path and local CLI configuration on each security appliance.

Note The URL for contacting this AUS server is produced by concatenating the `Protocol`://*Username*:*Password*@*IP IP Address*(:*Port*)/*Path* provided in these dialog boxes. The Port is included only if you entered a port number other than the default 443.

With the exception of the title, the Edit Auto Update Server dialog box is identical to the Add Auto Update Server dialog box. The following descriptions apply to both.

Navigation Path

You can access the Add and Edit Auto Update Server dialog boxes from the AUS Page, page 52-1.

Field Reference

Table 52-2 Add and Edit Auto Update Server Dialog Boxes

Element	Description
Protocol	The protocol used to communicate with the AUS server; choose **http** or **https**. **Note** If https is selected as the protocol to communicate with the Auto Update server, the security appliance will use SSL. This requires the security appliance to have a DES, 3DES, or AES license.
IP Address	Enter the IP address or Select a Networks/Hosts object representing this AUS server.
Port	Enter the number of the port on which communications with the AUS server take place. Defaults to 80 if http is chosen as the Protocol, and to 443 if https is chosen. If you enter an arbitrary port number, be sure the AUS server is configured to use the same port.
Path	The path to AUS services on the server. The standard path is `autoupdate/AutoUpdateServlet`; change this to `admin/auto-update` only if the AUS server host is an ASA.
AUS Interface	Enter or Select the interface to use when polling the Auto Update server.
Verify Certificate	Select this option to require SSL verification from the AUS server. The certificate returned by the server will be checked against Certification Authority (CA) root certificates. This requires that the AUS Server and this device use the same Certification Authority.
Username	Enter a user name to be used for AUS authentication (optional).
Password	Enter the password to be used for AUS authentication (optional).
Confirm	Re-enter the password (optional).

DHCP Relay Page

Use the DHCP Relay page to configure DHCP relay services for security devices. Dynamic Host Configuration Protocol (DHCP) relay passes DHCP requests received on one interface to an external DHCP server located behind a different interface. To configure DHCP relay, you need to specify at least one DHCP relay server and then enable a DHCP relay agent on the interface receiving DHCP requests.

Note You cannot enable a DHCP relay agent on an interface where a DHCP relay server is configured. The DHCP relay agent works only with external DHCP servers; it will not forward DHCP requests to a security appliance interface configured as a DHCP server.

Beginning with Security Manager version 4.9, DHCP Relay IPv4 is supported for ASA cluster devices running the software version 9.4.0 or higher.

For ASA-SM 9.1.2+ devices, you can configure DHCP relay servers per-interface, so requests that enter a given interface are relayed only to servers specified for that interface. When a DHCP request enters an interface that does not have interface-specific servers configured, the ASA relays the request to all global servers. If the interface has interface-specific servers, then the global servers are not used. IPv6 is not supported for per-interface DHCP relay. For more information, see Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DHCP Relay** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > DHCP Relay** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 52-3 DHCP Relay Page

Element	Description
DHCP Relay Agent table	This table lists the interfaces on which DHCP relay is configured. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Relay Agent Configuration dialog box, while Edit Row opens the Edit DHCP Relay Agent Configuration dialog box. See Add and Edit DHCP Relay Agent Configuration Dialog Boxes, page 52-6 for more information.
DHCP Servers table	This table lists the global DHCP servers to which DHCP requests are relayed. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Relay Server Configuration dialog box, while Edit Row opens the Edit DHCP Relay Server Configuration dialog box. See Add and Edit DHCP Relay Server Configuration Dialog Boxes, page 52-7 for more information.

Table 52-3 DHCP Relay Page (continued)

Element	Description
Timeout (seconds)	Specify the amount of time, in seconds, allowed for DHCP address negotiation. Valid values range from 1 to 3600 seconds; the default value is 60 seconds.
Trust Info (Option 82)	Specifies that you want to trust all DHCP client interfaces. You can configure interfaces as trusted interfaces to preserve DHCP Option 82. **Note** You can also specify interfaces to trust individually. For more information, see Add/Edit Interface Dialog Box: Advanced Tab (ASA/PIX 7.0+), page 46-41. DHCP Option 82 is used by downstream switches and routers for DHCP snooping and IP Source Guard. Normally, if the ASA DHCP relay agent receives a DHCP packet with Option 82 already set, but the giaddr field (which specifies the DHCP relay agent address that is set by the relay agent before it forwards the packet to the server) is set to 0, then the ASA will drop that packet by default. You can now preserve Option 82 and forward the packet by identifying an interface as a trusted interface.

Add and Edit DHCP Relay Agent Configuration Dialog Boxes

Use the Add DHCP Relay Agent Configuration dialog box to configure and enable a DHCP relay agent on an interface. Use the Edit DHCP Relay Agent Configuration dialog box to update an existing interface relay agent.

Note You cannot enable a DHCP relay agent on an interface where a DHCP relay server is configured. The DHCP relay agent works only with external DHCP servers; it will not forward DHCP requests to a security appliance interface configured as a DHCP server.

The Add DHCP Relay Agent Configuration dialog box and the Edit DHCP Relay Agent Configuration dialog box are virtually identical; the following descriptions apply to both.

Navigation Path

You can access the Add and Edit DHCP Relay Agent Configuration dialog boxes from the DHCP Relay Page, page 52-5.

Related Topics

- Add and Edit DHCP Relay Server Configuration Dialog Boxes, page 52-7

Field Reference

Table 52-4 Add and Edit DHCP Relay Agent Configuration Dialog Boxes

Element	Description
Interface	Enter or Select the name of the interface on which you want to configure a DHCP relay agent.

Table 52-4 Add and Edit DHCP Relay Agent Configuration Dialog Boxes (continued)

Element	Description
Enable DHCP Relay	When checked, the DHCP relay is enabled on the specified interface.
Set Route	Check this box to configure the DHCP relay agent to modify the default router address in the information returned from the DHCP server. When this option is selected, the DHCP relay agent substitutes the address of the selected interface for the default router address in the information returned from the DHCP server.

Add and Edit DHCP Relay Server Configuration Dialog Boxes

Use the Add DHCP Relay Server Configuration dialog box to define a new DHCP relay server; use the Edit DHCP Relay Server Configuration dialog box to update existing server information. You can configure a maximum of 10 DHCPv4 relay servers in single mode and per context, global and interface-specific servers combined, with a maximum of 4 servers per interface.

Note PIX Firewalls running an OS earlier than 7.2 only support 4 DHCP relay servers.

The Add DHCP Relay Server Configuration dialog box and the Edit DHCP Relay Server Configuration dialog box are virtually identical; the following descriptions apply to both.

Navigation Path

You can access the Add and Edit DHCP Relay Server Configuration dialog boxes from the DHCP Relay Page, page 52-5.

Related Topics

- Add and Edit DHCP Relay Agent Configuration Dialog Boxes, page 52-6

Field Reference

Table 52-5 Add and Edit DHCP Relay Server Configuration Dialog Boxes

Element	Description
Server	Enter the IP address or Select a Networks/Hosts object representing the external DHCP server to which DHCP requests are forwarded.
Interface	Enter or Select the interface through which DHCP requests are forwarded to the external DHCP server.

DHCP Relay IPv6 Page

Use the DHCP Relay IPv6 page to configure DHCPv6 relay services for security devices. Dynamic Host Configuration Protocol v6 (DHCPv6) relay passes DHCPv6 requests received on one interface to an external DHCPv6 server located behind a different interface. To configure DHCPv6 relay, you need to specify at least one DHCPv6 relay server and then enable a DHCPv6 relay agent on the interface receiving DHCPv6 requests.

Note You cannot enable a DHCPv6 relay agent on an interface where a DHCPv6 relay server is configured. The DHCPv6 relay agent works only with external DHCPv6 servers; it will not forward DHCPv6 requests to a security appliance interface configured as a DHCPv6 server.
Beginning with Security Manager version 4.9, DHCP Relay IPv6 is supported for ASA cluster devices running the software version 9.4.0 or higher.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DHCP Relay IPv6** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > DHCP Relay IPv6** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Note Two new interface settings have been introduced for DHCPv6: "managed-config-flag" and "other-config-flag." For more information, refer to Configuring IPv6 Interfaces (ASA/FWSM), page 46-47.

Field Reference

Table 52-6 DHCP Relay IPv6 Page

Element	Description
DHCP Relay IPv6 Agent table	This table lists the interfaces on which DHCP relay IPv6 is configured. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Relay IPv6 Agent Configuration dialog box, while Edit Row opens the Edit DHCP Relay IPv6 Agent Configuration dialog box. See Add and Edit DHCP Relay IPv6 Agent Configuration Dialog Boxes, page 52-9 for more information.
DHCP Servers table	This table lists the interfaces on which DHCP relay IPv6 is configured. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Relay IPv6 Server Configuration dialog box, while Edit Row opens the Edit DHCP Relay IPv6 Server Configuration dialog box. See Add and Edit DHCP Relay IPv6 Server Configuration Dialog Boxes, page 52-9 for more information.
Timeout (seconds)	Specify the amount of time, in seconds, allowed for DHCPv6 address negotiation. Valid values range from 1 to 3600 seconds; the default value is 60 seconds.

Add and Edit DHCP Relay IPv6 Agent Configuration Dialog Boxes

Use the Add DHCP Relay IPv6 Agent Configuration dialog box to configure and enable a DHCPv6 relay agent on an interface. Use the Edit DHCP Relay IPv6 Agent Configuration dialog box to update an existing interface relay agent.

Note You cannot enable a DHCPv6 relay agent on an interface where a DHCPv6 relay server is configured. The DHCPv6 relay agent works only with external DHCPv6 servers; it will not forward DHCPv6 requests to a security appliance interface configured as a DHCPv6 server.

The Add DHCP Relay IPv6 Agent Configuration dialog box and the Edit DHCP Relay IPv6 Agent Configuration dialog box are virtually identical; the following descriptions apply to both.

Navigation Path

You can access the Add and Edit DHCP Relay IPv6 Agent Configuration dialog boxes from the DHCP Relay IPv6 Page, page 52-7.

Related Topics

- Add and Edit DHCP Relay IPv6 Server Configuration Dialog Boxes, page 52-9

Field Reference

Table 52-7 Add and Edit DHCP Relay IPv6 Agent Configuration Dialog Boxes

Element	Description
Interface	Enter or Select the name of the interface on which you want to configure a DHCPv6 relay agent.
Enable DHCPv6 Relay	When checked, the DHCPv6 relay is enabled on the specified interface.
Set Route	Check this box to configure the DHCPv6 relay agent to modify the default router address in the information returned from the DHCPv6 server. When this option is selected, the DHCPv6 relay agent substitutes the address of the selected interface for the default router address in the information returned from the DHCPv6 server.

Add and Edit DHCP Relay IPv6 Server Configuration Dialog Boxes

Use the Add DHCP Relay IPv6 Server Configuration dialog box to define a new DHCPv6 relay server; use the Edit DHCP Relay IPv6 Server Configuration dialog box to update existing server information. You can define up to ten DHCPv6 relay servers.

Note The Add DHCP Relay IPv6 Server Configuration dialog box and the Edit DHCP Relay IPv6 Server Configuration dialog box are virtually identical; the following descriptions apply to both.

Navigation Path

You can access the Add and Edit DHCP Relay IPv6 Server Configuration dialog boxes from the DHCP Relay IPv6 Page, page 52-7.

Related Topics

- Add and Edit DHCP Relay IPv6 Agent Configuration Dialog Boxes, page 52-9

Field Reference

Table 52-8 Add and Edit DHCP Relay IPv6 Server Configuration Dialog Boxes

Element	Description
Server	Enter the IP address or Select a Networks/Hosts object representing the external DHCPv6 server to which DHCPv6 requests are forwarded.
Interface	Enter or Select the interface through which DHCPv6 requests are forwarded to the external DHCPv6 server.

Configuring DHCP Servers

A Dynamic Host Configuration Protocol (DHCP) server provides network configuration parameters, such as IP addresses, to DHCP clients. The security appliance can provide DHCP server or DHCP relay services to DHCP clients attached to the security appliance interfaces. The DHCP server provides network configuration parameters directly to DHCP clients: DHCP relay passes DHCP requests received on one interface to an external DHCP server located behind a different interface. For more information about DHCP relay, see DHCP Relay Page, page 52-5.

Note The security appliance DHCP server does not support BOOTP requests.

In multiple-context mode, you cannot enable a DHCP server or DHCP relay on an interface that is used by more than one context.

You can configure a DHCP server on each interface of the security appliance, and each interface can have its own pool of addresses to draw from. However, the other DHCP settings, such as DNS servers, domain name, options, ping timeout, and WINS servers, are configured globally and used by the DHCP server on all interfaces.

You cannot configure a DHCP client or DHCP relay services on an interface on which the DHCP server is enabled. Additionally, DHCP clients must be directly connected to the interface on which the server is enabled.

If your firewall is also acting as a DHCP client on the outside interface, you can enable auto-negotiated IP configuration. This allows the firewall to pass the DNS, WINS and domain name parameters it gets from the outside interface (as a DHCP client) to hosts on its inside network. Alternatively, you can manually specify the DNS, WINS and domain name parameters. If you specify those parameters manually and auto-configuration is on, your values take precedence over auto-configuration.

Use the DHCP Server Page, page 52-10 to manage DHCP server definitions.

DHCP Server Page

Use the DHCP Server page to configure global DHCP server and dynamic DNS (DDNS) update options, to set up a DHCP server on one or more device interfaces, and to configure advanced server options.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DHCP Server** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > DHCP Server** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Configuring DHCP Servers, page 52-10

Field Reference

Table 52-9 DHCP Server Page

Element	Description
Ping Timeout	Specify the amount of time, in milliseconds, that the firewall device waits to time out a DHCP ping attempt. To avoid address conflicts, firewall devices send two ICMP ping packets to an address before assigning that address to a DHCP client. Valid values range from 10 to 10000 milliseconds.
Lease Length	Specify the amount of time, in seconds, that the client can use its allocated IP address before the lease expires. Valid values range from 300 to 1048575 seconds. The default value is 3600 seconds (1 hour).
Enable auto-configuration (PIX and ASA only)	Select this option to enable DHCP auto configuration. DHCP auto configuration causes the DHCP server to provide DHCP clients with DNS server, domain name, and WINS server information obtained from a DHCP client running on the specified interface. If any of the information obtained through auto configuration is also specified manually, the manually specified information takes precedence over the discovered information.
Interface	If Enable auto-configuration is checked, this field is available. Enter or Select the interface running the DHCP client that supplies the DNS, WINS, and domain name parameters.
Define settings (optional)	
Domain Name	Specify the DNS domain name for DHCP clients. Enter a valid DNS domain name; for example, `example.com`.
Primary DNS Server	Enter the IP address or Select a Networks/Hosts object representing the primary DNS server for a DHCP client.
Primary WINS Server	Enter the IP address or Select a Networks/Hosts object representing the primary WINS server for a DHCP client.
Secondary DNS Server	Enter the IP address or Select a Networks/Hosts object representing the alternate DNS server for a DHCP client.
Secondary WINS Server	Enter the IP address or Select a Networks/Hosts object representing the alternate WINS server for a DHCP client.

Table 52-9 DHCP Server Page (continued)

Element	Description
Dynamic DNS Update	
Enable Dynamic DNS Update	Check this box to define global DDNS update options: • Select the type of resource-record updating: **PTR Record** only, or **A Record and PTR Record**. • You also can select Override DHCP Client Request. If selected, DHCP server updates override any updates requested by DHCP clients. These options are available only on ASA/PIX 7.2 and higher.
DHCP Server Interface Configuration table	
Interface table	This table lists device interfaces on which a DHCP server, DDNS updating, or both are configured. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Server Interface Configuration dialog box, while Edit Row opens the Edit DHCP Server Interface Configuration dialog box. See Add and Edit DHCP Server Interface Configuration Dialog Boxes, page 52-12 for more information.
Advanced Options	
Advanced button	Opens the Add/Edit DHCP Server Advanced Configuration Dialog Box, page 52-13.

Add and Edit DHCP Server Interface Configuration Dialog Boxes

Use these dialog boxes to enable DHCP and specify a DHCP address pool for a specified interface, and to enable dynamic DNS (DDNS) updating on the interface.

Note Other than the titles, the two dialog boxes are identical.

Navigation Path

You can access the Add DHCP Server Interface Configuration and Edit DHCP Server Interface Configuration dialog boxes from the DHCP Server Page, page 52-10.

Related Topics

- Configuring DHCP Servers, page 52-10

Field Reference

Table 52-10 Add/Edit DHCP Server Interface Configuration Dialog Boxes

Element	Description
Interface	Identifies the interface on which you are configuring a DHCP server. Enter an interface name, or select an interface object.

Table 52-10 Add/Edit DHCP Server Interface Configuration Dialog Boxes (continued)

Element	Description
DHCP Address Pool	Enter an IP address or a range of addresses, separated by a hyphen, that the DHCP server will use when assigning IP addresses. The beginning and ending addresses in the range must be in the same subnet, and the beginning address cannot be greater than the ending address.
Enable DHCP Server	Check this box to enable a DHCP server on this interface.
Enable Dynamic DNS Update	Check this box to enable DDNS updating by this DHCP server. Specify the record(s) to be updated: • **PTR Record** only • **A Record and PTR Record** You also can select **Override DHCP Client Request**. If selected, DHCP server updates override any updates requested by DHCP clients.

Add/Edit DHCP Server Advanced Configuration Dialog Box

The Add/Edit DHCP Server Advanced Configuration dialog box lets you manage DHCP options configured for the DHCP server. These options provide additional information to DHCP clients. For example, DHCP option 150 and DHCP option 66 provide TFTP server information to Cisco IP Phones and Cisco IOS routers.

Navigation Path

You can access the Add/Edit DHCP Server Advanced Configuration dialog box by clicking the Advanced button on the DHCP Server Page, page 52-10.

Related Topics

- Configuring DHCP Servers, page 52-10

Field Reference

Table 52-11 Add/Edit DHCP Server Advanced Configuration Dialog Box

Element	Description
Options table	This table lists configured DHCP server options. Use the Add Row, Edit Row, and Delete Row buttons to manage these entries. The Add Row button opens the Add DHCP Server Interface Configuration dialog box, while Edit Row opens the Edit DHCP Server Interface Configuration dialog box. See Add/Edit DHCP Server Option Dialog Box, page 52-13 for more information.

Add/Edit DHCP Server Option Dialog Box

The Add and Edit DHCP Server Option dialog boxes let you configure DHCP server option parameters, to provide additional information to DHCP clients. For example, DHCP option 150 and DHCP option 66 provide TFTP server information to Cisco IP Phones and Cisco IOS routers.

Navigation Path

You can access the Add and Edit DHCP Server Option dialog boxes from the Add/Edit DHCP Server Advanced Configuration Dialog Box, page 52-13.

Related Topics

- Configuring DHCP Servers, page 52-10
- DHCP Server Page, page 52-10

Field Reference

Table 52-12 Add/Edit DHCP Server Option Dialog Box

Element	Description
Option Code	Choose an option from the list of available option codes. All DHCP options (options 1 through 255) are supported except 1, 12, 50-54, 58-59, 61, 67, and 82. Detailed information about DHCP option codes is available on cisco.com: DHCP Options Reference.
Type	Choose the type of information the option returns to the DHCP client: • **IP** – Choosing this type specifies that one or two IP addresses are returned to the DHCP client. Provide up to two IP addresses. • **ASCII** – Choosing this type specifies that an ASCII value is returned to the DHCP client. Provide the ASCII character string, which cannot include spaces. • **HEX** – Choosing this type specifies that an hexadecimal value is returned to the DHCP client. Provide the HEX string with an even number of digits and no spaces; you do not need to use a 0x prefix.

DNS Page

Use the DNS page to configure DNS server groups. The firewall device uses these DNS servers to resolve fully-qualified domain names (host names) to IP addresses for SSL VPN, certificates, and FQDN network/host objects used in identity-aware firewall policies. Other features that define server names (such as AAA) do not support DNS resolution—you must enter the IP address or manually resolve the name to an IP address.

Tip The DefaultDNS server group is predefined on the ASA and is used for FQDN network/host object resolution. If you use FQDN objects, ensure that you configure DNS servers for this group; otherwise, the names cannot be resolved. To enhance security, ensure that you specify DNS servers that are trusted and that are preferably inside your network. For more information, see Requirements for Identity-Aware Firewall Policies, page 13-3.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DNS** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > DNS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add DNS Server Dialog Box, page 52-17

Field Reference

Table 52-13 DNS Page

Element	Description
DNS Server Groups table	This table lists the currently defined DNS server groups. Use the Add Row, Edit Row and Delete Row buttons below the table to manage these group entries. The Add Row button opens the Add DNS Server Group dialog box, and the Edit Row button opens the Edit DNS Server Group dialog box; except for the titles these dialog boxes are identical. See Add DNS Server Group Dialog Box, page 52-16 for more information.
DNS Lookup Interfaces	Lists the interfaces on which you want to enable DNS lookup. Enter or Select one or more interfaces or interface roles.
Enable DNS Guard (ASA/PIX 7.0(5), 7.2(x) and 8.x only)	Check this box to enable DNS Guard for the selected device or shared policy. DNS Guard tears down the DNS session associated with a DNS query as soon as the DNS reply is forwarded by the security appliance. DNS Guard also monitors the message exchange to ensure that the ID of the DNS reply matches the ID of the DNS query. This command is effective only on interfaces for which DNS inspection is disabled. When DNS inspection is enabled, the DNS Guard function is always performed. **Note** In releases prior to 7.0(5), the DNS Guard functions are always enabled regardless of the configuration of DNS inspection.

Table 52-13 DNS Page (continued)

Element	Description
DefaultDNS Server Group (ASA 8.4(2)+)	Additional settings that apply to the DefaultDNS server group only. These settings are used when resolving FQDN network/host objects to IP addresses. • **Poll Timer**—The time, in minutes, of the polling cycle used to resolve FQDN network/host objects to IP addresses. FQDN objects are resolved only if they are used in a firewall policy. The timer determines the maximum time between resolutions; the DNS entry's time-to-live (TTL) value is also used to determine when to update to IP address resolution, so individual FQDNs might be resolved more frequently than the polling cycle. The default is 240 (four hours). The range is 1 to 65535 minutes. • **Expire Entry Timer**—The number of minutes after a DNS entry expires (that is, the TTL has passed) that the entry is removed from the DNS lookup table. Removing an entry requires that the table be recompiled, so frequent removals can increase the processing load on the device. Because some DNS entries can have very short TTL (as short as three seconds), you can use this setting to virtually extend the TTL. The default is 1 minute (that is, the entry is removed one minute after the TTL has passed). The range is 1 to 65535 minutes.

Add DNS Server Group Dialog Box

Use the Add DNS Server Group dialog box to define the DNS servers and settings for a DNS server group, used by security devices to resolve server names to IP addresses in policies that support name resolution.

Note With the exception of its title, the Edit DNS Server Group dialog box is identical to this one, and the following descriptions apply to both.

Navigation Path

You can access the Add DNS Server Group and Edit DNS Server Group dialog boxes from the DNS Page, page 52-14.

Field Reference

Table 52-14 Add/Edit DNS Server Group Dialog Boxes

Element	Description
Name	Provide a name for the group of DNS servers. **Tip** The name DefaultDNS is predefined on the ASA and includes the servers used for policies that do not allow the selection of a specific group, such as for FQDN network/host object resolution.

Table 52-14 Add/Edit DNS Server Group Dialog Boxes (continued)

Element	Description
DNS Servers	Lists the DNS servers in this group. You can specify up to six servers to which DNS requests can be forwarded. The security appliance tries each DNS server in top-to-bottom order until it receives a response. **Note** You also must specify at least one interface on which DNS is enabled in the DNS Lookup section of the DNS Page, page 52-14. Use the buttons next to this list to manage the entries; from the top down, they are: • Add a DNS server to the list; opens the Add DNS Server Dialog Box, page 52-17. • Delete the currently selected DNS server entry from the list. • Move the currently selected entry up one row. • Move the currently selected entry down one row.
Timeout	Specify the number of seconds, from 1 to 30, to wait before trying the next DNS server; the default is 2 seconds. Each time the security device retries the list of servers, this timeout doubles.
Retries	Specify the number of times, from 0 to 10, to retry the list of DNS servers when the security device does not receive a response.
Domain Name	Optionally, specify a valid DNS domain name for the server; for example, **dnsexample.com**.

Add DNS Server Dialog Box

Use the Add DNS Server dialog box to add a DNS server to the DNS servers list in the Add DNS Server Group or Edit DNS Server Group dialog boxes.

Navigation Path

You can access the Add DNS Server dialog box from the Add DNS Server Group or Edit DNS Server Group dialog boxes. For more information about these dialog boxes, see Add DNS Server Group Dialog Box, page 52-16.

Related Topics

- DNS Page, page 52-14

Field Reference

Table 52-15 Add DNS Server Dialog Box

Element	Description
DNS Server	The IP address, or the host network/host object that defines the address, of the DNS server. Enter the address or click **Select** to select the network/host object from a list or to create a new object.

Table 52-15 Add DNS Server Dialog Box (continued)

Element	Description
Interface (ASA 9.5(1) or higher)	Click select to choose an interface. The Interface selector dialog box lists only Interface Roles and not Physical Interfaces. Therefore you must add a Physical Interface to the Interface Role before selecting the Source Interface. There is no default value for the interface. This feature is available in Security Manager version 4.9 and later for devices running ASA version 9.5(1) or higher.

Configuring DDNS

Dynamic DNS (DDNS) provides IP-address and domain-name mapping updates so hosts can find each other even though their DHCP-assigned IP addresses may change frequently. Also, beginning with the version 7.2(3), Cisco security appliances can generate DDNS updates. The DDNS page is where you configure this feature.

The DDNS mappings are maintained on the DHCP server in two types of resource records (RRs): the address or *A* records contain the name-to-IP-address mappings, while the pointer or *PTR* records map addresses to host names.

By automatically recording the association between assigned addresses and host names at defined intervals, DDNS allows frequently changing address-host name associations to be updated frequently. Mobile hosts, for example, can then move freely on a network without user or administrator intervention.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DDNS** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > DDNS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 52-16 DDNS Page

Element	Description
Dynamic DNS Interface Settings	This table lists currently defined DDNS interface-update methods. Use the Add Row, Edit Row, and Delete Row buttons below the table to manage these methods; the Add Row and Edit Row buttons open the Add/Edit DDNS Interface Rule Dialog Box, page 52-19.
DHCP Client requests DHCP Server to update records	The global setting on the appliance for DHCP client update requests. This option enables the client to send DDNS updates via the DHCP server, and specifies what is updated: the PTR resource record, both the A and PTR resource records, or neither. Choose **Not Selected**, **Only PTR Record**, **Both A and PTR Record**, or **No Update**.
DHCP Client ID Interface	Specify the interface(s) for global DHCP client update requests: enter an interface name or IP address, or Select an interface object.

Table 52-16 DDNS Page (continued)

Element	Description
Enable DHCP Client Broadcast	Select this option to allow DHCP clients on the device to broadcast DDNS updates. Available on ASA/PIX 7.2(3)+ devices only.

Add/Edit DDNS Interface Rule Dialog Box

Use the Add/Edit DDNS Interface Rule dialog box to manage rules for dynamic DNS updates. These rules are defined on a per-interface basis.

Navigation Path

You access the Add/Edit DDNS Interface Rule dialog box from the Configuring DDNS, page 52-18.

Related Topics

- DDNS Update Methods Dialog Box, page 52-19
- Add/Edit DDNS Update Methods Dialog Box, page 52-20

Field Reference

Table 52-17 Add/Edit DDNS Interface Rule Dialog Box

Element	Description
Interface	Enter or Select the name of the interface on which DDNS is to be configured. **Note** DHCP must be enabled on the specified interface.
Method Name	Choose a previously defined method for DDNS update, or choose **Add/Edit Update Method** to define a new method; the DDNS Update Methods Dialog Box, page 52-19 dialog box opens.
Hostname	Enter the name of the DDNS server host to which updates will be sent.
DHCP Client requests DHCP Server to update records	The setting on the interface for DHCP client update requests; specifies whether the DHCP server updates the PTR resource record, both the A and PTR records, or neither. Choose **Not Selected**, **Only PTR Record**, **Both A and PTR Record**, or **No Update**. Any choice other than Not Selected overrides the global setting on the Configuring DDNS, page 52-18.

DDNS Update Methods Dialog Box

Use the DDNS Update Methods dialog box to manage methods for dynamic DNS updates. Each defined method specifies an update interval and the resource record(s) to be updated.

Navigation Path

You access the DDNS Update Methods dialog box by choosing **Add/Edit Update Method** from the Method Name drop-down list in the Add/Edit DDNS Interface Rule Dialog Box, page 52-19.

Related Topics

- Configuring DDNS, page 52-18

Field Reference

Table 52-18 DDNS Update Methods Dialog Box

Element	Description
Update Methods	This table lists the currently defined update methods. Use the buttons below the table to manage these entries.
Add Row button	Opens the Add/Edit DDNS Update Methods Dialog Box, page 52-20 where you can define a new update method.
Edit Row button	Opens the Add/Edit DDNS Update Methods Dialog Box, page 52-20, where you can edit the method currently selected in the table.
Delete Row button	Deletes the method currently selected in the Update Methods table; confirmation may be required.

Add/Edit DDNS Update Methods Dialog Box

Use the Add/Edit DDNS Update Methods dialog box to define or edit a DDNS update method; currently defined methods are listed in the DDNS Update Methods Dialog Box, page 52-19.

Navigation Path

You access the Add/Edit DDNS Update Methods dialog box by clicking the Add Row or the Edit Row buttons in the DDNS Update Methods Dialog Box, page 52-19.

Related Topics

- Configuring DDNS, page 52-18

Field Reference

Table 52-19 Add/Edit DDNS Update Methods Dialog Box

Element	Description
Method Name	Provide an identifier for this method.
Update Interval	Specify how often records are to be updated for this method: provide a number of days, hours, minutes, and seconds. Note that while zero is the default value for hours, minutes and seconds, there is no default Day value: you must enter a number for Day.
Update Records	Specify the resource record(s) to be updated: select **Not Defined**, **A Records**, or **Both A and PTR Records**. Selecting A Records or Both A and PTR Records overrides the setting in the Add/Edit DDNS Interface Rule Dialog Box, page 52-19.

NTP Page

Network Time Protocol (NTP) is used to implement a hierarchical system of servers that provide precisely synchronized timing for network systems. This kind of accuracy is required for time-sensitive operations, such as validating Certificate Revocation Lists (CRLs), which include a precise time stamp. You can configure multiple NTP servers. The security device chooses the server with the lowest stratum—a measure of how reliable the data is.

Note This page is not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

Use the NTP page to enable NTP and manage the NTP servers used to dynamically set the time on a security device.

Note Time derived from an NTP server overrides any time set manually on the Clock page.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > NTP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > NTP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 52-20 NTP Page

Element	Description
Enable NTP Authentication	Enables or disables authentication with an NTP server. Disabling authentication does not alter the list of configured servers. If you enable authentication, the security appliance only communicates with an NTP server if it uses the correct trusted key in the packets. The security appliance also uses an authentication key to synchronize with the NTP server.
NTP Server Table	Lists the currently configured NTP servers. Use the Add Row, Edit Row and Delete Row buttons to manage this list; the Add Row and Edit Row buttons open the NTP Server Configuration Dialog Box, page 52-21.

NTP Server Configuration Dialog Box

Use the NTP Server Configuration dialog box to add or edit an NTP server definition.

Navigation Path

You can access the NTP Server Configuration dialog box from the NTP Page, page 52-21.

Note The NTP page is not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

Field Reference

Table 52-21 NTP Server Configuration Dialog Box

Element	Description
IP Address	Enter or Select the IP address of the NTP server.
Preferred	If checked, this NTP server is the preferred server when multiple servers are similarly accurate. NTP uses an algorithm to determine which server is the most accurate and synchronizes to that one. If multiple servers are of similar accuracy, then this option specifies which of those servers to use. However, if a server is significantly more accurate than the preferred one, the security appliance uses the more accurate server. For example, the security appliance uses a server of stratum 2 over a server of stratum 3 that is preferred. We recommend that you configure an NTP server as preferred only when multiple servers are likely to have the same stratum.
Interface	Enter or Select the interface used for NTP traffic, if you want to override the default interface in to the routing table.
Authentication Key The parameters in this section set the authentication key attributes if you want to use MD5 authentication for communicating with the NTP server.	
Key Number	Enter the ID for this authentication key. The NTP server packets must also use this key ID. If you previously configured a key ID for another server, you can select it in the list; otherwise, type a number between 1 and 4294967295.
Trusted	Sets this key as a trusted key. You must select this option for authentication to work.
Key Value	Enter the authentication key as a string up to 32 characters in length.
Confirm	Re-enter the authentication key to verify it is correct.

SMTP Server Page

Use the SMTP Server page to specify the IP address of an SMTP server and optionally, the IP address of a backup server, to which e-mail alerts and notifications are sent in response to specific events.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > SMTP Server** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > SMTP Server** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 52-22 SMTP Server Page

Element	Description
Primary Server IP Address	Enter or Select the IP address of the SMTP server.
Secondary Server IP Address	Enter or Select the IP address of a back-up SMTP server.

TFTP Server Page

The Trivial File Transfer Protocol (TFTP) is a simple client/server file transfer protocol described in RFC783 and RFC1350 Rev. 2. You can use the TFTP Server page to configure the security appliance as a TFTP client so it can transfer a copy of its running configuration file to a TFTP server. In this way, you can back up and propagate configuration files to multiple security appliances. Only one server is supported.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > TFTP Server** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > Server Access > TFTP Server** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 52-23 TFTP Server Page

Element	Description
Interface	Enter or Select the name of interface on which the TFTP server is accessed.
IP Address	Enter or Select the IP address of the TFTP server.
Directory	Enter the path on the TFTP server, beginning with a forward slash (/) and ending in the file name, to which the configuration files will be written (for example, **/tftpboot/asa/config3**). **Note** The path must begin with a forward slash (/).

Configuring FXOS Server Access Settings on Firepower 2100 Series Devices

The FXOS Server Access section contains pages for configuring FXOS server access on Firepower 2100 Series devices; FXOS Server Access is under Device Admin in the Device or Policy selector.

Note FXOS Server Access is not applicable for devices other than Firepower 2100 Series devices.

The Firepower 2100 Series devices supported by ASA and Cisco Security Manager are:

- Cisco FPR-2110 Adaptive Security Appliance
- Cisco FPR-2120 Adaptive Security Appliance
- Cisco FPR-2130 Adaptive Security Appliance
- Cisco FPR-2140 Adaptive Security Appliance

This chapter contains the following topics:

- HTTPS Page, page 53-1
- SSH Page, page 53-2
- SNMP Page, page 53-4

HTTPS Page

The HTTPS page allows you to configure the device to access the FXOS server through HTTPS. When you deploy configurations with this protocol, Cisco Security Manager encrypts the configuration file before sending it to the device.

Navigation Path

- (Device view) Select **Platform > Device Admin > FXOS Server Access > HTTPS** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > FXOS Server Access > HTTPS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add and Edit HTTPS Dialog Boxes, page 53-2

Field Reference

Table 53-1 HTTPS Page

Element	Description
Action	The permit action allows configuring the Firepower 2100 Series device with IP address and port. It supports IPv4 and IPv6 addresses.
Interface	The name of the device interface for which the HTTPS is configured. HTTPS cannot be configured on a Bridge Groups (BG) interface.
IP Address	The IP address of the device. It can be an IPv4 or an IPv6 address.
Port	The port on which communications with the FXOS server take place.

Add and Edit HTTPS Dialog Boxes

Use the Add HTTPS Configuration dialog box to create the HTTPS rules. The security appliance will automatically poll this server for image and configuration updates.

The Edit HTTPS Configuration dialog box is identical to the Add HTTPS Configuration dialog box. The following descriptions apply to both.

Navigation Path

You can access the Add and Edit HTTPS Configuration dialog boxes from the HTTPS Page, page 53-1.

Field Reference

Table 53-2 Add and Edit HTTPS Configuration Dialog Boxes

Element	Description
Action	Select Permit.
Interface	Click Select and choose the interface. The Bridge Groups (BG) interface cannot be configured with HTTPS.
IP Address	Click Select and choose the IP address of the device that can access the FXOS server. It can be an IPv4 or an IPv6 address.
Port	This value defaults to 3443 when you save the page. You can also enter the port on which communications with the FXOS server should take place.

SSH Page

Use the Secure Shell page to configure port that permit FXOS server access to a Firepower 2100 Series device using the SSH protocol. The rules permit SSH access to a specific IP address and netmask.

Navigation Path

- (Device view) Select **Platform > Device Admin > FXOS Server Access > SSH** from the Device Policy selector.

- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > FXOS Server Access > SSH** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add and Edit SSH Dialog Boxes, page 53-3

Field Reference

Table 53-3 SSH Page

Element	Description
Action	The permit action allows configuring the Firepower 2100 Series device with IP address and port to access the FXOS server. It supports IPv4 and IPv6 addresses.
Interface	The name of the device interface for which the SSH is configured. SSH cannot be configured on a Bridge Groups (BG) interface.
IP Address	The IP address of the device. It can be an IPv4 or an IPv6 address.
Port	The port on which communications with the FXOS server take place.

Add and Edit SSH Dialog Boxes

Use the Add SSH Configuration dialog box to create the SSH rules. The security appliance will automatically poll this server for image and configuration updates.

The Edit SSH Configuration dialog box is identical to the Add SSH Configuration dialog box. The following descriptions apply to both.

Navigation Path

You can access the Add and Edit SSH Configuration dialog boxes from the SSH Page, page 53-2.

Field Reference

Table 53-4 Add and Edit SSH Configuration Dialog Boxes

Element	Description
Action	Select Permit.
Interface	Click Select and choose the interface. The Bridge Groups (BG) interface cannot be configured with SSH.
IP Address	Click Select and choose the IP address of the device that can access the FXOS server. It can be an IPv4 or an IPv6 address.
Port	This value defaults to 3022 when you save the page. You can also enter the port on which communications with the FXOS server should take place.

SNMP Page

SNMP is an application layer protocol that facilitates the exchange of management information between network devices. You can use the SNMP page to configure the Firepower 2100 Series devices for monitoring by SNMP.

Navigation Path

- (Device view) Select **Platform > Device Admin > FXOS Server Access > SNMP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Device Admin > FXOS Server Access > SNMP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add and Edit SNMP Dialog Boxes, page 53-4

Field Reference

Table 53-5 SNMP Page

Element	Description
Action	The permit action allows configuring the Firepower 2100 Series device with IP address and port to access the FXOS server. It supports IPv4 and IPv6 addresses.
Interface	The name of the device interface for which the SNMP is configured. SNMP cannot be configured on a Bridge Groups (BG) interface.
IP Address	The IP address of the device. It can be an IPv4 or an IPv6 address.
Port	The port on which communications with the FXOS server take place.

Add and Edit SNMP Dialog Boxes

Use the Add SNMP Configuration dialog box to create the SNMP rules. The security appliance will automatically poll this server for image and configuration updates.

The Edit SNMP Configuration dialog box is identical to the Add SNMP Configuration dialog box. The following descriptions apply to both.

Navigation Path

You can access the Add and Edit SNMP Configuration dialog boxes from the SNMP Page, page 53-4.

Field Reference

Table 53-6 Add and Edit SNMP Configuration Dialog Boxes

Element	Description
Action	Select Permit.
Interface	Click Select and choose the interface. The Bridge Groups (BG) interface cannot be configured with SNMP.

Table 53-6 Add and Edit SNMP Configuration Dialog Boxes (continued)

Element	Description
IP Address	Click Select and choose the IP address of the device that can access the FXOS server. It can be an IPv4 or an IPv6 address.
Port	This value defaults to 3161 when you save the page. You can also enter the port on which communications with the FXOS server should take place.

Configuring Logging Policies on Firewall Devices

The Logging feature lets you enable and manage NetFlow "collectors," and enable system logging, set up logging parameters, configure event lists (syslog filters), apply the filters to a destination, set up syslog messages, configure syslog servers, and specify e-mail notification parameters.

After you enable logging and set up the logging parameters using the Logging Setup page, the Event Lists page lets you configure filters (for a set of syslogs) which can be sent to a logging destination. The Logging Filters page lets you specify a logging destination for the syslogs to be sent. Finally, the Syslog and E-Mail pages configure syslog and e-mail setup.

This chapter contains the following topics:

- NetFlow Page, page 54-1
- Embedded Event Manager, page 54-3
- E-Mail Setup Page, page 54-8
- Event Lists Page, page 54-9
- Logging Filters Page, page 54-12
- Configuring Logging Setup, page 54-14
- Configuring Rate Limit Levels, page 54-17
- Configuring Syslog Server Setup, page 54-20
- Defining Syslog Servers, page 54-26

NetFlow Page

A device configured for NetFlow data export captures flow-based traffic statistics on the device. This information is periodically transmitted from the device to a NetFlow collection server, in the form of User Datagram Protocol (UDP) datagrams.

The NetFlow page lets you enable NetFlow export on the selected device, and define and manage NetFlow "collectors" to which collected flow information is transmitted.

Navigation Path

- (Device view) Select **Platform > Logging > NetFlow** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > NetFlow** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Using Rules Tables, page 12-8
- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49

Field Reference

Table 54-1 NetFlow Page

Element	Description
Enable Flow Export	If checked, NetFlow data export is enabled.
Template Export Interval	Interval (in minutes) between transmissions of flow information to the collectors. The value can be from one to 3600 minutes; the default is 30.
Active Refresh Interval	For active connections, specifies the time interval between flow-update events in minutes. Valid values are from 1 to 60 minutes. The default value is 1 minute.
Delay Flow Create	Delays the sending of a flow-create event by the specified number of seconds. The value can be from one to 180 seconds. If no value is entered, there is no delay, and the flow-create event is exported as soon as the flow is created. If the flow is torn down before the configured delay, the flow-create event is not sent; an extended flow teardown event is sent instead.
Collectors table	Lists the currently defined NetFlow collectors. Use the Add Row, Edit Row and Delete Row buttons below the table to manage these entries. The Add Row and Edit Row buttons open the Add and Edit Collector Dialog Boxes (NetFlow), page 54-2. **Note** Cisco Security Manager does not allow duplicate netflow collectors for ASA 9.6(4) to 9.7.0, and 9.8(2) and above devices. Change the current configuration or remove the duplicate or overlapping configuration (Platform> Logging > Netflow) for the device.

Add and Edit Collector Dialog Boxes (NetFlow)

Use the Add Collector and Edit Collector dialog boxes to define and edit NetFlow "collectors." Except for the title, the two dialog boxes are identical; the following information applies to both.

Navigation Path

You can open the Add and Edit Collector dialog boxes from the NetFlow Page, page 54-1.

Field Reference

Table 54-2 Add and Edit Collector Dialog Boxes

Element	Description
Interface	Enter or Select the name of the device interface through which the collector is contacted.

Table 54-2 Add and Edit Collector Dialog Boxes (continued)

Element	Description
Collector	Enter the IP address or the network name of the server to which NetFlow packets will be sent. You also can Select a Networks/Hosts object.
UDP Port	Specify the UDP port on the specified Collector to which NetFlow packets will be sent. Values can range from 1 to 65535; the default is 2055.

Embedded Event Manager

The Embedded Event Manager (EEM) enables you to debug problems and provides general purpose logging for troubleshooting. There are two components: events to which the EEM responds or listens, and event manager applets that define actions as well as the events to which the EEM responds. You may configure multiple event manager applets to respond to different events and perform different actions.

Note Embedded Event Manager is supported on ASA 9.2(1)+ only.

Supported Events

The EEM supports the following events:

- Syslog—The ASA uses syslog message IDs to identify syslog messages that trigger an event manager applet. You may configure multiple syslog events, but the syslog message IDs may not overlap within a single event manager applet.
- Timers—You may use timers to trigger events. You may configure each timer only once for each event manager applet. Each event manager applet may have up to three timers. The three types of timers are the following:
  - Watchdog (periodic) timers trigger an event manager applet after the specified time period following the completion of the applet's actions and restart automatically.
  - Countdown (one-shot) timers trigger an event manager applet once after the specified time period and do not restart unless they are removed, then re-added.
  - Absolute (once-a-day) timers cause an event to occur once a day at a specified time, and restart automatically. The time-of-day format is in hh:mm:ss.

  You may configure only one timer event of each type for each event manager applet.
- None—The none event is triggered when you run an event manager applet manually.
- Crash—The crash event is triggered when the ASA crashes. Regardless of the value of the output command, the action commands are directed to the crashinfo file. The output is generated before the show tech command.

Note Be careful when using a range of Syslog IDs and when using timers. Incorrect configuration can cause an ASA loop and prevent the applet from executing normally.

Configuring Actions

When an event manager applet is triggered, the actions on the event manager applet are performed. Each action has a number that is used to specify the sequence of the actions. The sequence number must be unique within an event manager applet. You may configure multiple actions for an event manager applet. The commands are typical CLI commands, such as **show blocks**.

Configuring Output Destinations

You may send the output of the action CLI commands to one of three locations:

- None, which is the default and discards the output
- Console, which sends the output to the ASA console
- File, which sends the output to a file. The following four file options are available:
  - new—creates a new, uniquely named file each time that an event manager applet is invoked.
  - overwrite—overwrites a specified file each time that an event manager applet is invoked.
  - append—appends to a specified file each time that an event manager applet is invoked. If the file does not yet exist, it is created.
  - rotate—creates a set of uniquely named files that are rotated each time that an event manager applet is invoked.

Guidelines and Limitations

- Supported in single mode only. Not supported in multiple context mode.
- Supported in routed and transparent firewall modes.
- EEM will be enabled irrespective of whether logging functionality is enabled on the device or not.
- The EEM functionality on the ASA only contains a subset of the EEM functionality found on Cisco routers.
- During a crash, the state of the ASA is generally unknown. Some commands may not be safe to run during this condition.
- The name of an event manager applet may not contain spaces.
- You cannot modify the None event and Crashinfo event parameters.
- Performance may be affected because syslog messages are sent to the EEM for processing.
- The default output is none for each event manager applet. To change this setting, you must enter a different output value.
- You may have only one output option defined for each event manager applet.

The Embedded Event Manager table lists the currently defined event manager applets. Use the Add Row, Edit Row and Delete Row buttons below the table to manage these entries. The Add Row and Edit Row buttons open the Add and Edit Applet Dialog Boxes, page 54-5.

Navigation Path

- (Device view) Select **Platform > Logging > Embedded Event Manager** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Embedded Event Manager** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add and Edit Applet Dialog Boxes, page 54-5
- Table Columns and Column Heading Features, page 1-49

Add and Edit Applet Dialog Boxes

Use the Add Applet and Edit Applet dialog boxes to define and edit event manager applets. Except for the title, the two dialog boxes are identical; the following information applies to both.

Navigation Path

You can open the Add and Edit Applet dialog boxes from the Embedded Event Manager, page 54-3.

Field Reference

Table 54-3 Add and Edit Applet Dialog Boxes

Element	Description
Name	Enter a unique name for the event manager applet. The name cannot contain spaces and must be less than 32 characters.
Description	Enter a description for the event manager applet. The description may be up to 256 characters long.
Configuration Tab	
Crashinfo	When selected, the event manager applet is triggered when the ASA crashes. Regardless of the value of the Output field, the action commands are directed to the crashinfo file. The output is generated before the show tech command. **Note** The state of the ASA is generally unknown when it crashes. Some CLI commands may not be safe to run during this condition.
None	When selected, you can trigger the event manager applet manually. **Note** Manual triggering of the EEM applet is not supported in Cisco Security Manager. To manually trigger an applet, you must use a FlexConfig. See Chapter 7, "Managing FlexConfigs" for more information.
Syslog table	The Syslog table lists the currently defined syslog message IDs for the selected applet. Use the Add Row, Edit Row and Delete Row buttons below the table to manage these entries. The Add Row and Edit Row buttons open the Add and Edit Syslog Configuration Dialog Boxes, page 54-7.
Absolute	Configure an absolute (once-a-day) timer event Absolute timers cause an event to occur once a day at a specified time, and restart automatically. Use the fields provided to enter the time of day in hours, minutes, and seconds. The time range is from 00:00:00 (midnight) to 23:59:59.
Countdown	Configures a countdown (one-shot) timer event. Countdown timers trigger an event manager applet once after the specified time period and do not restart unless they are removed, then re-added. Enter the time period in seconds. The number of seconds may range from 1-604800.

Table 54-3 Add and Edit Applet Dialog Boxes (continued)

Element	Description
Watchdog	Configures a watchdog (periodic) timer event. Watchdog timers trigger an event manager applet after the specified time period following the completion of the applet's actions and restart automatically. Enter the time period in seconds. The number of seconds may range from 1-604800.
Output	To configure specific destinations for sending output from an action, choose one of the available output destination options: • none—(default) discards the output. • console—sends the output to the ASA console. • file—sends the output to a file. Select the file option in the Action list.
Action	The following four file options are available: • new—creates a new, uniquely named file each time that an event manager applet is invoked. The filename has the format of eem-applet-timestamp.log, in which applet is the name of the event manager applet and timestamp is a dated timestamp in the format of YYYYMMDD-hhmmss. • overwrite—overwrites a specified file each time that an event manager applet is invoked. Specify the file details using the File Location and File Name fields. • append—appends to a specified file each time that an event manager applet is invoked. If the file does not yet exist, it is created. Specify the file details using the File Location and File Name fields. • rotate—creates a set of uniquely named files that are rotated each time that an event manager applet is invoked. Specify the number of files to be rotated in the File Count field (valid values range from 2 - 100). When a new file is to be written, the oldest file is deleted, and all subsequent files are renumbered before the first file is written. The newest file is indicated by 0, and the oldest file is indicated by the highest number. The filename format is eem-applet-x.log, in which applet is the name of the applet, and x is the file number.
File Location	Specifies the location of the output file. The location may also use FTP, TFTP, and SMB targeted files.
File Name	Specifies the filename of the output file.
File Count	Specify the number of files to be rotated when "rotate" is the selected Action. When a new file is to be written, the oldest file is deleted, and all subsequent files are renumbered before the first file is written. The newest file is indicated by 0, and the oldest file is indicated by the highest number. Valid values for the rotate value range from 2 - 100. The filename format is eem-applet-x.log, in which applet is the name of the applet, and x is the file number.

Table 54-3 Add and Edit Applet Dialog Boxes (continued)

Element	Description
Action Tab	
Action table	The Action table lists the currently defined actions for the selected applet. Use the Add Row, Edit Row and Delete Row buttons below the table to manage these entries. The Add Row and Edit Row buttons open the Add and Edit Action Configuration Dialog Boxes, page 54-7.

Add and Edit Syslog Configuration Dialog Boxes

Use the Add Syslog Configuration and Edit Syslog Configuration dialog boxes to configure the syslog message IDs for an event manager applet. Except for the title, the two dialog boxes are identical; the following information applies to both.

Navigation Path

You can open the Add and Edit Syslog Configuration dialog boxes from the Add and Edit Applet Dialog Boxes, page 54-5.

Field Reference

Table 54-4 Add and Edit Syslog Configuration Dialog Boxes

Element	Description
ID	Enter a single syslog message or a range of syslog messages. If a syslog message occurs that matches the specified individual syslog message or range of syslog messages, an event manager applet is triggered. **Note** Syslog message IDs may not be entered twice or overlap within a single event manager applet.
Occurs	(Optional) In the occurrences field, enter the number of times that the syslog message must occur for an event manager applet to be invoked. The default is 1 occurrence every 0 seconds. Valid values are from 1 - 4294967295.
Period	(Optional) In the period field, enter the number of seconds within which the syslog messages must occur to invoke the action. This value limits how frequently an event manager applet is invoked to at most once in the configured period. Valid values are from 0 - 604800. A value of 0 means that no period is defined.

Add and Edit Action Configuration Dialog Boxes

Use the Add Action Configuration and Edit Action Configuration dialog boxes to configure the actions for an event manager applet. Except for the title, the two dialog boxes are identical; the following information applies to both.

Navigation Path

You can open the Add and Edit Action Configuration dialog boxes from the Add and Edit Applet Dialog Boxes, page 54-5.

Field Reference

Table 54-5 Add and Edit Action Configuration Dialog Boxes

Element	Description
Ordinal ID	Enter the unique sequence number in the Ordinal ID field. Valid sequence numbers range from 0 - 4294967295. When adding an action configuration, the Ordinal ID will default to one greater than the highest Ordinal ID used.
CLI	Enter the CLI command in the CLI field. The command runs in global configuration mode as a user with privilege level 15 (the highest). The command may not accept any input, because it is disabled.

E-Mail Setup Page

The E-Mail Setup page (PIX 7.0/ASA Only) lets you set up a source e-mail address, as well as a list of recipients for specified syslog messages to be sent as e-mails. You can filter the syslog messages sent to a destination e-mail address by severity. The table shows which entries have been set up.

The syslog severity filter used for the destination e-mail address will be the higher of the severity selected in this section and the global filter set for all e-mail recipients in the Logging Filters page.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > E-Mail Setup** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > E-Mail Setup** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 54, “Configuring Logging Policies on Firewall Devices”

Field Reference

Table 54-6 E-Mail Setup Page

Element	Description
Source Email Address	Enter the email address to be used as the source address when syslogs are sent as emails.
Destination Address table	Lists the currently defined email recipients of syslog messages. Use the Add Row, Edit Row and Delete Row buttons below the table to manage this list; the Add Row and Edit Row buttons open the Add/Edit Email Recipient Dialog Box, page 54-8.

Add/Edit Email Recipient Dialog Box

The Add/Edit Email Recipient dialog box lets you configure a destination address to be sent emails containing syslog messages; you can limit the messages sent according to severity.

The syslog severity filter used for the destination email address will be the higher of the severity selected in this section and the global filter set for all email recipients on the Logging Filters Page, page 54-12.

Navigation Path

You can access the Add/Edit Email Recipient dialog box from the E-Mail Setup Page, page 54-8.

Field Reference

Table 54-7 Add/Edit Email Recipient Dialog Box

Element	Description
Destination Email Address	Enter the recipient email address for the chosen type of syslog messages.
Syslog Severity list	Choose the severity of the syslogs to be emailed to this recipient; messages of the chosen severity and higher are sent. Message severity levels are described in Logging Levels, page 54-24.

Event Lists Page

The Event Lists page (PIX 7.0+/ASA only) lets you define a set of syslog message filters for logging. After you enable logging and set up global logging parameters on the Logging Setup page, use this page to configure event lists used to filter syslog messages sent to different logging destinations. (The Logging Filters Page, page 54-12 lets you specify logging destinations for event lists.)

Use the Add Row, Edit Row and Delete Row buttons below the Event Lists table to manage the entries. Add Row and Edit Row open the Add/Edit Event List Dialog Box, page 54-10.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Event Lists** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Event Lists** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Logging Setup Page, page 54-15
- Chapter 54, "Configuring Logging Policies on Firewall Devices"

Message Classes and Associated Message ID Numbers

The following table lists the message classes and the range of message IDs in each class.

Table 54-8 Message Classes and Associated Message ID Numbers

Class	Definition	Message ID Numbers
auth	User Authentication	109, 113
bridge	Transparent Firewall	110, 220
ca	PKI Certification Authority	717
config	Command interface	111, 112, 208, 308
e-mail	E-mail Proxy	719
ha	Failover (High Availability)	101, 102, 103, 104, 210, 311, 709

Table 54-8 Message Classes and Associated Message ID Numbers (continued)

Class	Definition	Message ID Numbers
ids	Intrusion Detection System	400, 401, 415
ip	IP Stack	209, 215, 313, 317, 408
np	Network Processor	319
ospf	OSPF Routing	318, 409, 503, 613
rip	RIP Routing	107, 312
rm	Resource Manager	321
session	User Session	106, 108, 201, 202, 204, 302, 303, 304, 305, 314, 405, 406, 407, 500, 502, 607, 608, 609, 616, 620, 703, 710
snmp	SNMP	212
sys	System	199, 211, 214, 216, 306, 307, 315, 414, 604, 605, 606, 610, 612, 614, 615,701, 711
vpdn	PPTP and L2TP Sessions	213, 403, 603
vpn	IKE and IPsec	316, 320, 402, 404, 501, 602, 702, 713, 714, 715
vpnc	VPN Client	611
vpnfo	VPN Failover	720
vpnlb	VPN Load Balancing	718
webvpn	Web-based VPN	716

Add/Edit Event List Dialog Box

The Add/Edit Event List dialog box lets you create or edit an event list, and specify which syslog messages to include in the event list filter.

You can use the following criteria to define an event list:

- Class and Severity
- Message ID

Class represents specific types of related syslog messages. For example, the class `auth` represents all syslog messages related to user authentication.

Severity classifies syslogs based on the relative importance of the event in the normal functioning of the network. The highest severity is Emergency, which means the resource is no longer available. The lowest severity is Debugging, which provides detailed information about every network event.

The message ID is a numeric value that uniquely identifies each individual message. You can specify a single message ID, or a range of IDs, in an event list.

Navigation Path

You can access the Add/Edit Event List dialog box from the Event Lists Page, page 54-9.

Related Topics

- Chapter 54, “Configuring Logging Policies on Firewall Devices”

Field Reference

Table 54-9 Add/Edit Event List Dialog Box

Element	Description
Event List Name	Enter a name that uniquely identifies this event list.
Event Class/Severity Filters	This table lists the event class and severity level filters defined for this event list. Use the Add Row, Edit Row and Delete Row buttons below this table to manage the entries. Add Row and Edit Row open the Add/Edit Syslog Class Dialog Box, page 54-11.
Message ID Filters	This table list the message ID filters defined for this event list. Use the Add Row, Edit Row and Delete Row buttons below this table to manage the entries. Add Row and Edit Row open the Add/Edit Syslog Message ID Filter Dialog Box, page 54-11.

Add/Edit Syslog Class Dialog Box

The Add/Edit Syslog Class dialog box lets you specify an event class and a related severity level as an event list filter.

Class represents specific types of related syslog messages, so you do not have to select the syslogs individually. For example, the class `auth` represents all syslog messages related to user authentication.

Severity classifies syslogs based on the relative importance of the event in the normal functioning of the network. The highest severity is Emergency, which means the resource is no longer available. The lowest severity is Debugging, which provides detailed information about every network event.

Navigation Path

You access the Add/Edit Syslog Class dialog box from the Add/Edit Event List Dialog Box, page 54-10.

Related Topics

- Add/Edit Syslog Message ID Filter Dialog Box, page 54-11
- Event Lists Page, page 54-9

Field Reference

Table 54-10 Add/Edit Syslog Class Dialog Box

Element	Description
Event Class	Choose the desired event class. Event classes are described in Table 54-8 on page 54-9.
Severity	Choose the desired message severity level. Severity levels are described in Logging Levels, page 54-24.

Add/Edit Syslog Message ID Filter Dialog Box

The Add/Edit Syslog Message ID Filter dialog box lets you specify a syslog message ID, or a range of IDs, as an the event list filter.

Navigation Path

You can access the Add/Edit Syslog Message ID Filter dialog box from the Add/Edit Event List Dialog Box, page 54-10.

Related Topics

- Add/Edit Syslog Class Dialog Box, page 54-11
- Event Lists Page, page 54-9

Field Reference

Message IDs – Enter a syslog message ID, or a range of IDs. Use a hyphen to specify a range; for example, `101001-101010`. Message IDs must be between 100000 and 999999.

Message IDs and their corresponding messages are listed in the System Log Message guides for the appropriate product. You can access these guides from cisco.com:

PIX Firewall

- http://www.cisco.com/en/US/products/sw/secursw/ps2120/products_system_message_guides_list.html

ASA

- http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html

FWSM

- http://www.cisco.com/en/US/products/hw/modules/ps2706/ps4452/tsd_products_support_model_home.html

Logging Filters Page

The Logging Filters page lets you configure a logging destination for event lists (syslog filters) that have been configured using the Event Lists page, or for only the syslog messages that you specify using the Edit Logging Filters page. Syslog messages from specific or all event classes can be selected using the Edit Logging Filters page.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Logging Filters** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Logging Filters** from the Policy Type selector. Right-click **Logging Filters** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Chapter 54, “Configuring Logging Policies on Firewall Devices”
- Edit Logging Filters Dialog Box, page 54-13

Field Reference

Table 54-11 Logging Filters Page

Element	Description
Logging Destination	Lists the name of the logging destination to which messages matching this filter are sent. Logging destinations are as follows: • **Internal Buffer**. Messages matching this filter are published to the internal buffer of the security appliance. • **Console**. Messages matching this filter are published to any console port connections. • **Telnet Sessions**. Messages matching this filter are published to any Telnet sessions connected to the security appliance. • **Syslog Servers**. Messages matching this filter are published to any syslog servers specified on the **Platform > Logging > Syslog Servers page**. • **E-Mail**. Messages matching this filter are published to any recipients specified on the **Platform > Logging > E-mail Setup (PIX7.0/ASA Only)** page. • **SNMP Trap**. Messages matching this filter are published to any SNMP management stations specified on the **Platform > Device Admin > Device Access > SNMP** page. • **ASDM**. Messages matching this filter are published to any ASDM sessions.
Syslogs From All Event Classes	Lists the severity on which to filter, the event list to use, or whether logging is disabled from all event classes. Event classes are described in Message Classes and Associated Message ID Numbers, page 54-9.
Syslogs From Specific Event Classes	Lists event class and severity set up as the filter. Event classes are described in Message Classes and Associated Message ID Numbers, page 54-9. Severity levels are described in Logging Levels, page 54-24.

Edit Logging Filters Dialog Box

The Edit Logging Filters dialog box lets you edit filters for a logging destination. Syslogs can be configured from all or specific event classes, or disabled for a specific logging destination.

Navigation Path

You can access the Edit Logging Filters dialog box from the Logging Filters page. For more information about the Logging Filters page, see Logging Filters Page, page 54-12.

Related Topics

- Chapter 54, “Configuring Logging Policies on Firewall Devices”
- Logging Filters Page, page 54-12

Field Reference

Table 54-12 Edit Logging Filters Dialog Box

Element	Description
Logging Destination list	Specifies the logging destination for this filter: • **Internal Buffer**. Messages matching this filter are published to the internal buffer of the security appliance. • **Console**. Messages matching this filter are published to any console port connections. • **Telnet Sessions**. Messages matching this filter are published to any Telnet sessions connected to the security appliance. • **Syslog Servers**. Messages matching this filter are published to any syslog servers specified on the **Platform > Logging > Syslog Servers page**. • **E-Mail**. Messages matching this filter are published to any recipients specified on the **Platform > Logging > E-mail Setup (PIX7.0/ASA Only)** page. • **SNMP Trap**. Messages matching this filter are published to any SNMP management stations specified on the **Platform > Device Admin > Device Access > SNMP** page. • **ASDM**. Messages matching this filter are published to any ASDM sessions.
Syslog from All Event Classes	
Filter on severity option	Filters on the severity of the logging messages.
Filter on severity list	Specifies the level of logging messages on which to filter.
Use event list option	Specifies to use an event list.
Use event list	Specifies the event list to use. Event lists are defined on the Event Lists Page, page 54-9.
Disable logging option	Disables all logging to the selected destination.
Syslog from Specific Event Classes (PIX7.0)	
Event Class	Specifies the event class and severity. Event classes include one or all available items. Event classes are described in Table 54-8 on page 54-9.
Severity	Specifies the level of logging messages. Severity levels are described in Table 54-18 on page 54-24.

Configuring Logging Setup

The Logging Setup page lets you enable system logging on the security appliance and configure other logging options. These options include enabling logging on the security appliance and failover unit, specifying the base log format and detail, and logging to longer-term storage devices, FTP server or Flash, before purging the internal buffer.

Related Topics

- Logging Setup Page, page 54-15

Step 1 Select **Platform > Logging > Syslog > Logging Setup** to display the Logging Setup page.

Step 2 Check **Enable Logging**.

This option enables logging on the security appliance.

Step 3 To enable logging on the failover unit paired with this security appliance, select the **Enable logging on the standby failover unit** check box.

Step 4 To enable EMBLEM format, or to send debug messages as part of the syslog messages, select the corresponding check boxes.

If you enable EMBLEM, you must use the UDP protocol to publish syslog messages. It is not compatible with TCP.

Step 5 To write the internal buffer data to an FTP server for future processing prior to clearing the buffer, do the following:

a. Check **FTP Server Buffer wrap**.

b. Enter the IP address of the FTP server in the **IP Address** field.

c. Enter the user name of the account used to log into the FTP server in the **User Name** field.

d. Enter the path in the **Path** field, relative to the FTP root, where the file should be stored.

e. Enter and confirm the password used to authenticate the user name.

Step 6 To write the internal buffer data to Flash for future processing prior to clearing the buffer, do the following:

a. Check **Flash**.

b. Specify the maximum amount of memory to allocate to the storage of internal buffer data.

c. Specify the minimum memory that should remain free on the Flash drive. If this minimum value cannot be retained while writing out the data from the internal buffer, the messages will be pruned to meet the space requirements.

Step 7 To specify the maximum queue size maintained on the appliance for viewing by an ASDM client, enter that value in the **Message Queue Size (Messages)** field.

Logging Setup Page

The Logging Setup page lets you enable system logging on the security appliance and configure other logging options.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Logging Setup** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Logging Setup** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 54, “Configuring Logging Policies on Firewall Devices”

Field Reference

Table 54-13 Logging Setup Page

Element	Description
Enable Logging	Turns on logging for the main security appliance.
Enable Logging on the Failover Standby Unit	Turns on logging for the standby security appliance, if available.
Send syslogs in EMBLEM format (PIX7.x+, ASA, FWSM 3.x+)	Enables EMBLEM format logging for every logging destination. If you enable EMBLEM, you must use the UDP protocol to publish syslog messages; EMBLEM is not compatible with TCP. **Note** This setting is not compatible with CS-MARS.
Send debug messages as syslogs (PIX7.x+, ASA, FWSM 3.x+)	Redirects all the debug trace output to the syslog. The syslog message does not appear in the console if this option is enabled. Therefore, to see debug messages, you must enable logging at the console and configure it as the destination for the debug syslog message number and logging level. The syslog message number used is *711011*. Default logging level for this syslog is *debug*.
Memory Size of Internal Buffer (bytes)	Specify the size of the internal buffer to which syslogs is saved if the logging buffer is enabled. When the buffer fills up, it is overwritten. The default is 4096 bytes. The range is 4096 to 1048576.
Specify FTP Server Information (PIX7.x+, ASA, FWSM 3.x+)	
FTP Server Buffer Wrap	To save the buffer contents to the FTP server before it is overwritten, check this box and enter the necessary destination information in the following fields. To remove the FTP configuration, deselect this option.
IP Address	Enter the IP address of the FTP server.
User Name	Enter the user name to use when connecting to the FTP server.
Path	Enter the path, relative to the FTP root, where the buffer contents should be saved.
Password/Confirm	Enter and confirm the password used to authenticate the user name to the FTP server.
Specify flash size	
Flash	To save the buffer contents to the flash memory before it is overwritten, check this box. This option is only available in routed or transparent single mode.
Maximum flash to be used by logging (KB)	Specify the maximum space to be used in the flash memory for logging (in KB). This option is available only in routed or transparent single mode.
Minimum free space to be preserved (KB)	Specifies the minimum free space to be preserved in flash memory (in KB). This option is available only in routed or transparent single mode.
ASDM Logging (PIX7.x+, ASA, FWSM 3.x+)	
Message Queue Size	Specify the queue size for syslogs intended for viewing in ASDM.

Configuring Rate Limit Levels

The Rate Limit page lets you specify the maximum number of log messages of specific types (e.g., “alert” or “critical”), and messages with specific Syslog IDs, that can be generated within given periods of time. You can specify individual limits for each logging level, and each Syslog message ID. If the settings conflict, the Syslog message ID limits take precedence.

The Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19 is used to specify the maximum number of messages that can be generated for a particular Syslog message ID within a given period of time.

The Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18 is used to specify the maximum number of messages that can be generated for a particular Syslog logging level within a given period of time.

Related Topics

- Rate Limit Page, page 54-18

Follow these steps to manage rate limits for message logging:

Step 1 Access the Rate Limit page by doing one of the following:

- (Device view) Select **Platform > Logging > Syslog > Rate Limit** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Rate Limit** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new policy.

Step 2 Add, edit and delete rate limits for Syslog logging levels:

- To specify the maximum number of messages that can be generated within a given period of time for particular logging level, click the **Add Row** button under the Rate Limits for Syslog Logging Levels table to open the Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18. Choose a logging level and define a rate limit.
- To edit the rate limit for a particular logging level, select the appropriate entry in the Rate Limits for Syslog Logging Levels table, and then click the **Edit Row** button under the table to open the Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18. Alter the rate limit as necessary.
- To delete a rate limit entry from the Rate Limits for Syslog Logging Levels table, select it and then click the **Delete Row** button under the table. A confirmation dialog box may be displayed; click OK to delete the entry.

Step 3 Add, edit and delete limits for log messages according to message IDs:

- To specify the maximum number of messages that can be generated within a given period of time for particular message ID, click the **Add Row** button under the Individually Rate Limited Syslog Messages table to open the Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19. Choose a Syslog message ID and define a rate limit.
- To edit the rate limit for a particular Syslog message ID, select the appropriate entry in the Individually Rate Limited Syslog Messages table, and then click the **Edit Row** button under the table to open the Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19. Alter the rate limit as necessary.

- To delete a message limit entry from the Individually Rate Limited Syslog Messages table, select it and then click the **Delete Row** button under the table. A confirmation dialog box may be displayed; click OK to delete the entry.

Rate Limit Page

The Rate Limit page allows you to specify the maximum number of log messages of a particular type (for example, alert or critical) that should be generated within a given period of time. You can specify a limit for each logging level and Syslog message ID. If the settings differ, Syslog message ID limits take precedence.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Rate Limit** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Rate Limit** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new policy.

Related Topics

- Chapter 54, "Configuring Logging Policies on Firewall Devices"
- Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18
- Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19

Field Reference

Table 54-14 Rate Limit Page

Element	Description
Rate Limits for Syslog Logging Levels Table	
Logging Level	The Syslog logging level for which you are specifying a rate limit.
No. of Messages	Maximum number of messages of the specified type allowed in the specified time period.
Interval (seconds)	Number of seconds before the rate limit counter resets.
Individually Rate Limited Syslog Messages Table	
Syslog ID	Identification number of the Syslog message for which you are specifying a rate limit.
No. of Messages	Maximum number of messages with the specified ID allowed in the specified time period.
Interval (seconds)	Number of seconds before the rate limit counter resets.

Add/Edit Rate Limit for Syslog Logging Levels Dialog Box

Using the Add/Edit Rate Limit for Syslog Logging Levels dialog box, you can specify the maximum number of log messages for particular log level that should be generated within a given period of time. You can specify a limit for each logging level or syslog message ID (see Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19). If the settings differ, the rate limited syslog message-level settings override rate limit logging level settings.

Navigation Path

You can access the Add/Edit Rate Limit for Syslog Logging Levels dialog box from the Rate Limit page. For more information, see Rate Limit Page, page 54-18.

Related Topics

- Chapter 54, "Configuring Logging Policies on Firewall Devices"
- Rate Limit Page, page 54-18
- Add/Edit Rate Limited Syslog Message Dialog Box, page 54-19

Field Reference

Table 54-15 Add/Edit Rate Limit for Syslog Logging Levels Dialog Box

Element	Description
Logging Level	The syslog logging level for which you are specifying the rate limit.
Number of Messages	Maximum number of messages of the specified type allowed in the specified time period.
Interval (Seconds)	Number of seconds before the rate limit counter resets.

Add/Edit Rate Limited Syslog Message Dialog Box

Using the Add/Edit Rate Limited Syslog Message dialog box you can specify the maximum number of log messages of a particular Syslog ID that can be generated within a given period of time. You can specify a limit for each syslog message ID or logging level (see Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18). If the settings differ, the rate limited syslog message-level settings override rate limit logging level settings.

Navigation Path

You can access the Add/Edit Rate Limited Syslog Message dialog box from the Rate Limit page. For more information, see Rate Limit Page, page 54-18.

Related Topics

- Chapter 54, "Configuring Logging Policies on Firewall Devices"
- Rate Limit Page, page 54-18
- Add/Edit Rate Limit for Syslog Logging Levels Dialog Box, page 54-18

Field Reference

Table 54-16 Add/Edit Rate Limited Syslog Message Dialog Box

Element	Description
Syslog ID	Identification number of the syslog message for which you are specifying a rate limit.
Number of Messages	Maximum number of messages with the specified ID allowed in the specified time period.
Interval (Seconds)	Number of seconds before the rate limit counter resets.

Configuring Syslog Server Setup

You can configure general syslog server settings to set the facility code to be included in syslog messages that are sent to syslog servers, specify whether a timestamp is included in each message, specify the device ID to include in messages, view and modify the severity levels for messages, and disable the generation of specific messages.

Related Topics

- Defining Syslog Servers, page 54-26

Step 1 Do one of the following:

- (Device view) Select **Platform > Logging > Syslog > Server Setup** to open the Server Setup Page, page 54-21.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Server Setup** from the Policy Type selector. Select an existing policy or create a new one.

Step 2 Change the basic message configuration as required:

- If your syslog server expects a different facility than the default, select the required facility in the **Facility** list.
- If you want to include the date and time a message was generated in the message, select **Enable Timestamp on Each Syslog Message**.
- If you want to add a device identifier to syslog messages (which is placed at the beginning of the message), select **Enable Syslog Device ID** and then select the type of ID:

  Note For an ASA cluster, each unit in the cluster generates its own syslog messages. You can configure logging so that each unit uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all units in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all units look as if they come from a single unit. If you configure logging to use the local-unit name that is assigned in the cluster bootstrap configuration as the device ID (Cluster ID option), syslog messages look as if they come from different units. You can also specify whether or not the interface IP address of the control unit should be used for all cluster devices.

  - **Interface**—To use the IP address of the specified interface, regardless of the interface through which the appliance sends the message. Click **Select** to select the interface or the interface role that identifies the interface. Interface roles must map to a single interface.

    For ASA clusters, to specify that the interface IP address of the Control Unit should be used for all cluster devices, select the corresponding option under the Interface Name field.
  - **User Defined ID**—To use a text string (up to 16 characters) of your choosing.
  - **Host Name**—To use the hostname of the device.
  - **Cluster ID**—To use the unique name in the boot configuration of an individual ASA unit in the cluster as the device ID.

Step 3 Use the Syslog Message table to alter the default settings for specific syslog messages. You need to configure rules in this table only if you want to change the default settings. You can change the severity assigned to a message, or you can suppress (disable) the generation of a message.

- To add a rule, click the **Add Row** button and fill in the Add/Edit Syslog Message Dialog Box, page 54-25.

  You select the message number whose configuration you want to change, and then select the new severity level, or select **Suppressed** to disable the generation of the message. Typically, you would not change the severity level and disable the message, but you can make changes to both fields if desired. Click **OK** to add the rule to the table.

  For a description of message severity levels, see Logging Levels, page 54-24.
- To edit a rule, select it and click the **Edit Row** button, make the desired changes, and click **OK**.
- To delete a rule, select it and click the **Delete Row** button.
- If you are using NetFlow, you can easily disable the generation of syslog messages that have NetFlow equivalents by clicking the **Disable NetFlow Equivalent Syslogs** button. This adds the messages to the table as suppressed messages. Note that if any of these syslog equivalents are already in the table, your existing rules are not overwritten.

Syslog Relay Configuration

In addition to events being received by the Cisco Security Manager server, they can be forwarded to a maximum of two external/remote controllers (syslog hosts). Syslog relay will forward the received messages to another syslog host using the UDP syslog protocol.

If you want the syslog messages that are forwarded from the Cisco Security Manager server to have the Cisco Security Manager server's IP address as the source IP address of the syslog message, you must enable it through CLI command:

1. Navigate to CSCOpx\MDC\logrelay and open the logrelay.properties file.
2. Set the values of ext1 and ext2 to false like this:

```
## Source Preservation
#logrelay.dp.txring.ext0.preserve.source=true
logrelay.dp.txring.ext1.preserve.source=false
logrelay.dp.txring.ext2.preserve.source=false
```

Note By default the value is true for all collectors, by setting ext1 and ext2 as false, Cisco Security Manager will send the sylog messages with Cisco Security Manager IP. This modification can be done only for remote collectors and not for the local collector (ext0).

Server Setup Page

The Server Setup page allows you to set the facility code to be included in syslog messages that are sent to syslog servers, specify whether a timestamp is included in each message, specify the device ID to include in messages, view and modify the severity levels for messages, and disable the generation of specific messages.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Server Setup** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Server Setup** from the Policy Type selector. Select an existing policy or create a new one.

Related Topics

- Configuring Syslog Server Setup, page 54-20
- Defining Syslog Servers, page 54-26
- Chapter 54, "Configuring Logging Policies on Firewall Devices"
- Logging Levels, page 54-24

Field Reference

Table 54-17 Server Setup Page

Element	Description
Facility	The syslog facility code that the appliance includes in messages destined for syslog servers. The default is LOCAL4(20), which is what most UNIX systems expect. You can select a facility between LOCAL0(16) and LOCAL7(23). Syslog facility is useful when you have a central syslog monitoring system that needs to distinguish among the various network devices that generate syslog data streams. Because your network devices share the eight available facilities, you might need to change this value.
Enable Timestamp on Each Syslog Message	Whether to include the date and time a message was generated in syslog messages. The default is to not include time stamps.

Table 54-17 Server Setup Page (continued)

Element	Description
Enable Syslog Device ID	Whether to configure a device ID in non-EMBLEM-format syslog messages. If you select this option, select one of the following to use as the device ID, which is place at the start of all syslog messages: **Note** For an ASA cluster, each unit in the cluster generates its own syslog messages. You can configure logging so that each unit uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all units in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all units look as if they come from a single unit. If you configure logging to use the local-unit name that is assigned in the cluster bootstrap configuration as the device ID (Cluster ID option), syslog messages look as if they come from different units. You can also specify whether or not the interface IP address of the Control Unit should be used for all cluster devices. • Interface—The IP address of the selected interface. Enter the name of the interface or click **Select** to select it from a list (or to select an interface role that specifies the interface). Messages include the IP address of the interface specified, regardless of which interface the adaptive security appliance uses to send the log data to the external server. If you select an interface role, that role must map to a single interface on the device. For ASA clusters, to specify that the interface IP address of the Control Unit should be used for all cluster devices, select the corresponding option under the Interface Name field. • User Defined ID—A text string you define as the device ID. This string can be up to 16 characters, but cannot contain any of the following special characters: & ’ " < > ? • Host Name—The hostname of the security appliance. • Cluster ID—Use the unique name in the boot configuration of an individual ASA unit in the cluster as the device ID.

Table 54-17 Server Setup Page (continued)

Element	Description
Syslog Message table	Use this table to enable or disable the generation of specific syslog messages, or to change the severity level of a message. If you do not want to constrict which message types are generated, or change any message severity levels, you do not need to configure anything in this table. The table shows the messages you have configured with the message level and whether generation is suppressed ("true" in the table). • To add a rule, click the **Add Row** button and fill in the Add/Edit Syslog Message Dialog Box, page 54-25. • To edit a rule, select it and click the **Edit Row** button. • To delete a rule, select it and click the **Delete Row** button.
Disable/Enable NetFlow Equivalent Syslogs	If you are using NetFlow logging, you might want to disable the generation of syslog messages that duplicate NetFlow messages. If you click the Disable button, these duplicate syslog messages are added to the Syslog Message table as suppressed messages, and the button is renamed Enable NetFlow Equivalent Syslogs. Clicking the Enable button removes the duplicate syslog messages from the table, meaning that they will no longer be suppressed, and the device will start sending them again. However, if you manually edited any message that was added to the list by the Disable button, the Enable button does not remove them.

Logging Levels

The following table describes logging levels.

Table 54-18 Logging Levels

Logging Level	Type	Description
0	Emergency	System unusable. Generates messages that identify system instabilities.
1	Alerts	Immediate action needed. Generates messages that identify system integrity issues that require immediate administrative action.
2	Critical	Critical condition. Generates messages that identify critical system issues.
3	Errors	Error condition. Generates messages that identify system errors during operation.
4	Warnings	Warning condition. Generates messages that identify system warnings. For example, device might be configured incorrectly.
5	Notifications	Normal but significant condition. Generates messages that identify normal operations that are typically considered significant events.
6	Information	Informational only. Generates messages that identify system information that is typical of day-to-day activity, such as network session records.

Table 54-18 Logging Levels (continued)

Logging Level	Type	Description
7	Debugging	Generates syslog messages that assist you in debugging. Also generates logs that identify the commands issued during FTP sessions and the URLs requested during HTTP sessions. Includes all emergency, alert, critical, error, warning, notification, and information messages.
-	Disabled	No logging.

Add/Edit Syslog Message Dialog Box

The Add/Edit Syslog Message dialog box lets you modify the logging level or suppression setting for a syslog message.

Navigation Path

You can access the Add/Edit Syslog Message dialog box from the Server Setup Page, page 54-21.

Related Topics

- Chapter 54, "Configuring Logging Policies on Firewall Devices"

Field Reference

Table 54-19 Add/Edit Syslog Message Dialog Box

Element	Description
Syslog ID list	The message log ID of the message whose severity level or suppression setting you want to alter. These values and their corresponding messages are identified in the System Log Message guides for the appropriate product: **PIX Firewall** http://www.cisco.com/en/US/products/sw/secursw/ps2120/products_system_message_guides_list.html **ASA** http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html **FWSM** http://www.cisco.com/en/US/products/hw/modules/ps2706/ps4452/tsd_products_support_model_home.html **Note** Starting from Cisco Security Manager 4.10, you can enter a syslog message in the **Syslog ID** field. Make sure that you enter a valid syslog ID corresponding to the device; else the deployment may fail.
Logging Level list	The logging level that you want to assign to the message. For logging levels and descriptions, see Logging Levels, page 54-24. Select (default) to use the default level assigned to the message.
Suppressed	Whether to suppress the generation of the syslog message. Suppressing a message disables its generation, so you will not see it in syslogs.

Table 54-19 Add/Edit Syslog Message Dialog Box (continued)

Element	Description
Disable Syslogs on Standby	Whether to block specific syslog messages from being generated on standby ASA devices. This feature is available from ASA version 9.4(1) and Security Manager supports this feature starting from version 4.9.

Defining Syslog Servers

The Syslog Servers page lets you specify the syslog servers to which the security appliance will send syslog messages. To make use of the syslog servers you define, you must enable logging using the Logging Setup page and set up the appropriate filters for destinations using the Logging Filters page.

Tip If you want to view events from an ASA device using Security Manager Event Viewer, ensure that you define the Security Manager server as a syslog server. Also, if you use CS-MARS or other applications to manage syslog events, include those servers in this policy.

By directing syslog records generated by a security appliance to a syslog server, you can process and study the records.

Before You Begin

Enable logging. See Configuring Logging Setup, page 54-14.

Related Topics

- Syslog Servers Page, page 54-27
- Add/Edit Syslog Server Dialog Box, page 54-28

Step 1 Select **Platform > Logging > Syslog > Syslog Servers** to display the Syslog Servers page.

Step 2 Do one of the following:

- To add a new syslog target, click the **Add Row** button.
- To edit an existing syslog target, select the check box for the row, then click the **Edit Row** button.

Step 3 Enter or select the interface name in the **Interface** field.

The list displays all interfaces defined at the current scope.

Step 4 Enter or select the IP address of the syslog server in the **IP Address** field.

Step 5 Determine whether to use UDP or TCP, then click the appropriate radio button under Protocol.

Step 6 Enter the port from which the security appliance sends either UDP or TCP syslog messages. The port must be the same port on which the syslog server listens.

- TCP—1470 (Default). TCP ports work only with a security appliance syslog server.
- UDP—514 (Default).

Step 7 To generate syslog messages using the EMBLEM format, select the **Log messages in Cisco EMBLEM format** check box.

To enable this option, you must select UDP protocol to publish messages to this syslog server.

Step 8 Click **OK**.

The definition appears in the Syslog Servers table.

Syslog Servers Page

The Syslog Servers page lets you specify the syslog servers to which the security appliance sends syslog messages. To make use of the syslog servers you define, you must enable logging using the Logging Setup page and set up the appropriate filters for destinations using the Logging Filters page.

Tip If you want to view events from an ASA device using Security Manager Event Viewer, ensure that you define the Security Manager server as a syslog server. Also, if you use CS-MARS or other applications to manage syslog events, include those servers in this policy.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog > Syslog Servers** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Logging > Syslog > Syslog Servers** from the Policy Type selector. Select an existing policy or create a new one.

Related Topics

- Defining Syslog Servers, page 54-26
- Chapter 54, "Configuring Logging Policies on Firewall Devices"

Field Reference

Table 54-20 Syslog Servers Page

Element	Description
Syslog Servers table	The syslog servers to which this device sends syslog messages. The table shows the device interface that publishes messages to the server, the server's IP address, syslog protocol and port number, and whether the messages are in Cisco EMBLEM syslog format. There is a limit of four syslog servers that can be set up per context. • To add a server, click the **Add Row** button and fill in the Add/Edit Syslog Server Dialog Box, page 54-28. • To edit a server, select it and click the **Edit Row** button. • To delete a server, select it and click the **Delete Row** button.
Queue Size	Specifies the size of the queue for storing syslog messages on the security appliance when syslog server is busy. Minimum is 1 message. Default is 512. Specify 0 to allow an unlimited number of messages to be queued (subject to available block memory).
Allow user traffic to pass when TCP syslog server is down	Whether to restrict all traffic if any syslog server that is using the TCP protocol is down.

Add/Edit Syslog Server Dialog Box

The Add/Edit Syslog Servers dialog box lets you add or edit the syslog servers to which the security appliance will send syslog messages. To make use of the syslog servers you define, you must enable logging using the Logging Setup page and set up the appropriate filters for destinations using the Logging Filters page.

Note There is a limit of four syslog servers that can be set up per context.

Navigation Path

You can access the Add Syslog Servers dialog box from the Syslog Servers page. For more information about the Syslog Servers page, see Syslog Servers Page, page 54-27.

Related Topics

- Defining Syslog Servers, page 54-26
- Chapter 54, “Configuring Logging Policies on Firewall Devices”

Field Reference

Table 54-21 Add/Edit Syslog Server Dialog Box

Element	Description
Interface	The interface used to communicate with the syslog server. Enter the name of the interface or interface role object, or click **Select** to select it from a list or to create a new object.
IP Address	The IP address of syslog server. Enter the IP address or the name of the network/host policy object that defines the address, or click **Select** to select the network/host object. **Note** Starting with Cisco Security Manager 4.13, IPv6 addresses are supported for the syslog server.
Protocol	The protocol used by syslog server, either TCP or UDP. UDP is the default. TCP ports work only with a security appliance syslog server. **Note** You must select UDP if you intend to use the EMBLEM format.
Port	The TCP or UDP port from which the security appliance sends syslog messages and on which the syslog server receives them. The default ports for each protocol are: • TCP—1470. • UDP—514. **Tip** If you are defining the Security Manager server as a syslog server, you can find the port number on the Security Manager Administration Event Management Page, page 11-27. **Note** During the installation or upgrade of Security Manager, the Common Services syslog service port is changed from 514 to 49514. Later, if Security Manager is uninstalled, the port is not reverted to 514.

Table 54-21 Add/Edit Syslog Server Dialog Box (continued)

Element	Description
Log messages in Cisco EMBLEM format (UDP only)	Whether to log messages in Cisco EMBLEM format. The syslog server must use UDP. **Note** If the syslog server is a Cisco Security MARS appliance, do not select this option. Cisco Security MARS does not process the EMBLEM format.
Reference Identity	Beginning with version 4.12, Security Manager enables you to select Reference Identity policy object name from the Policy Objects Selector. Reference Identity is enabled only if the Port is TCP and is disabled if the Port is UDP. For more information, see Reference Identities, page 49-13.

Configuring Multicast Policies on Firewall Devices

The Multicast section contains pages for defining IP multicast routing on security devices. Multicast routing is supported in single-context, routed mode only.

Enabling multicast routing enables IGMP and PIM on all interfaces by default. Internet Group Management Protocol (IGMP) is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. Protocol Independent Multicast (PIM) is used to maintain forwarding tables for multicast datagrams.

Note Only the UDP transport layer is supported for multicast routing.

This chapter contains the following topics:

- Enabling PIM and IGMP, page 55-1
- Configuring IGMP, page 55-2
- Configuring Multicast Routes, page 55-8
- Configuring Multicast Boundary Filters, page 55-9
- Configuring PIM, page 55-11

Enabling PIM and IGMP

The Enable PIM and IGMP page lets you enable or disable Internet Group Management Protocol (IGMP) and Protocol Independent Multicast (PIM) on all interfaces on the security appliance. IGMP is used to learn whether members of a group are present on directly attached subnets. Hosts join multicast groups by sending IGMP report messages. PIM is used to maintain forwarding tables to forward multicast datagrams.

When **Enable PIM and IGMP** is checked on this page, PIM and IGMP are enabled on all interfaces on the security appliance. Deselect the option to disable PIM and IGMP on all interfaces.

Note You can disable PIM and IGMP on a per-interface basis; see IGMP Page - Protocol Tab, page 55-3 and PIM Page - Protocol Tab, page 55-11 for more information.

Navigation Path

- (Device view) Select **Platform > Multicast > Enable PIM and IGMP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Multicast > Enable PIM and IGMP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Configuring IGMP, page 55-2
- Configuring Multicast Routes, page 55-8
- Configuring Multicast Boundary Filters, page 55-9
- Configuring PIM, page 55-11

Configuring IGMP

Internet Protocol hosts use IGMP to report their group memberships to directly connected multicast routers. Internet Group Management Protocol (IGMP) uses group-address (Class D) IP addresses.

Host group addresses can be in the range 224.0.0.0 to 239.255.255.255. The address 224.0.0.0 is never assigned to any group. The address 224.0.0.1 is assigned to all systems on a subnet. The address 224.0.0.2 is assigned to all routers on a subnet.

The IGMP page provides four tabbed panels, used to configure and manage IGMP in Security Manager:

- IGMP Page - Protocol Tab, page 55-3 – This panel displays interface-specific IGMP parameters; you can disable IGMP and change IGMP parameters.
- IGMP Page - Access Group Tab, page 55-5 – Lets you manage access groups that restrict the multicast sources allowed on an interface.
- IGMP Page - Static Group Tab, page 55-6 – Sometimes, hosts on a network may have a configuration that prevents them from answering IGMP queries; however, you still want multicast traffic to be forwarded to that network segment. There are two methods to pull multicast traffic down to a network segment:
  - Use the Join Group tab to configure the interface as a member of the multicast group. With this method, the security appliance accepts the multicast packets in addition to forwarding them to the specified interface.
  - Use the Static Group tab to configure the security appliance to be a statically connected member of a group. With this method, the security appliance does not accept the packets itself, but only forwards them. Therefore, this method allows fast switching. The outgoing interface appears in the IGMP cache, but itself is not a member of the multicast group.

  Use this tab to statically assign a multicast group to an interface, or change existing static group assignments.
- IGMP Page - Join Group Tab, page 55-7 – Use this tab to manage the multicast groups to which the security appliance belongs.

Note If you simply want to forward multicast packets for a specific group to an interface without the security appliance accepting those packets as part of the group, see IGMP Page - Static Group Tab, page 55-6.

Navigation Path

- (Device view) Select **Platform > Multicast > IGMP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Multicast > IGMP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

IGMP Page - Protocol Tab

Use the Protocol tab to configure IGMP parameters for an interface on the security appliance.

Navigation Path

You can access the Protocol tab from the IGMP page. For more information about the IGMP page, see Configuring IGMP, page 55-2.

Related Topics

- Configure IGMP Parameters Dialog Box, page 55-4
- Enabling PIM and IGMP, page 55-1
- Configuring PIM, page 55-11
- Configuring Multicast Routes, page 55-8

Field Reference

Table 55-1 Protocol Tab

Element	Description
Protocol Table	
Interface	The name of the interface to which the IGMP settings apply.
Enabled	Indicates whether IGMP is enabled on the interface.
Version	The version of IGMP enabled on the interface.
Query Interval	The interval, in seconds, at which the designated router sends IGMP host-query messages. Valid values range from 1 to 3600 seconds. The default value is 125 seconds.
Query Timeout	The period of time, in seconds, before the security appliance takes over querying the interface, after the previous appliance has stopped doing so. Valid values range from 60 to 300 seconds. The default value is 255 seconds.
Response Time	The maximum response time, in seconds, advertised in IGMP queries. If the security appliance does not receive any host reports within the designated response time, the IGMP group is pruned. Decreasing this value lets the security appliance prune groups faster. Valid values range from 1 to 12 seconds. The default value is 10 seconds. Changing this value is only valid only for IGMP Version 2.
Group Limit	The maximum number of hosts that can join on an interface. Valid values range from 1 to 500. The default value is 500.
Maximum Groups (PIX 6.3)	The maximum number of groups enabled for multicast. Valid values range from 0 to 2000.

Table 55-1 Protocol Tab (continued)

Element	Description
Forward Interface	The name of the interface to which the selected interface forwards IGMP host reports if IGMP forwarding is enabled.

Configure IGMP Parameters Dialog Box

Use the Configure IGMP Parameters dialog box to configure IGMP parameters for an interface on the security appliance.

Navigation Path

You can access the Configure IGMP Parameters dialog box from the IGMP Page - Protocol tab. For more information, see IGMP Page - Protocol Tab, page 55-3.

Related Topics

- IGMP Page - Protocol Tab, page 55-3
- Configuring IGMP, page 55-2

Field Reference

Table 55-2 Configure IGMP Parameters Dialog Box

Element	Description
Interface	The name of the interface to which the IGMP settings apply.
Forward Interface	The name of the interface to which IGMP host reports are forwarded if IGMP forwarding is enabled.
Version	The version of IGMP to enable on the interface. Choose 1 to enable IGMP Version 1, or 2 to enable IGMP Version 2. Some features require IGMP Version 2. By default, the security appliance uses IGMP Version 2.
Query Interval	The interval, in seconds, at which the designated router sends IGMP host-query messages. Valid values range from 1 to 3600 seconds. The default value is 125 seconds.
Response Time	The maximum response time, in seconds, advertised in IGMP queries. If the security appliance does not receive any host reports within the designated response time, the IGMP group is pruned. Decreasing this value lets the security appliance prune groups faster. Valid values range from 1 to 12 seconds. The default value is 10 seconds. Changing this value is valid only for IGMP Version 2.
Maximum Groups (PIX 6.3)	The maximum number of groups enabled for multicast. Valid values range from 0 to 2000.
PIX 7.x and ASA Only	
Enable IGMP	When selected, IGMP is enabled on the specified interface.
Group Limit	The maximum number of hosts that can join on an interface. Valid values range from 1 to 500. The default value is 500.

Table 55-2 Configure IGMP Parameters Dialog Box (continued)

Element	Description
Query Timeout	The period of time, in seconds, before the security appliance takes over querying the interface, after the previous appliance has stopped doing so. Valid values range from 60 to 300 seconds. The default value is 255 seconds.

IGMP Page - Access Group Tab

Use the Access Group tab to control the multicast groups that are allowed on an interface.

The table on this page lists all currently defined multicast access groups, showing for each, the name of the interface or interface role for which the group is defined, the group network(s), and whether this group is permitted or denied. For a detailed explanation of these fields, see Configure IGMP Access Group Parameters Dialog Box, page 55-5.

- To add a multicast access group to the table, click the Add Row button.
- To edit the settings for a group, select it and click the Edit Row button.
- To delete a group, select it and click the Delete Row button.

Navigation Path

You can access the Access Group tab from the Configuring IGMP, page 55-2.

Related Topics

- Enabling PIM and IGMP, page 55-1
- Configuring Multicast Routes, page 55-8

Configure IGMP Access Group Parameters Dialog Box

Use the Configure IGMP Access Group Parameters dialog box to add or modify an access group entry.

Navigation Path

You can access the Configure IGMP Access Group Parameters dialog box from the IGMP Page - Access Group Tab, page 55-5.

Related Topics

- IGMP Page - Access Group Tab, page 55-5
- Configuring IGMP, page 55-2

Field Reference

Table 55-3 Configure IGMP Access Group Parameters Dialog Box

Element	Description
Interface	Enter or Select the name of the interface to which the access group is assigned.

Table 55-3 Configure IGMP Access Group Parameters Dialog Box (continued)

Element	Description
Multicast Group Network	Enter or Select the multicast group address(es) assigned to the specified interface. You can provide one or more IP address/netmask entries, one or more Networks/Hosts objects, or a combination of both; separate the entries with commas. Group network addresses can range from 224.0.0.0 to 239.255.255.255.
Action	Choose **permit** if the multicast group is permitted on the interface. Choose **deny** if the multicast group is not permitted.

IGMP Page - Static Group Tab

Use the Static Group tab to statically assign a multicast group to an interface.

Navigation Path

You can access the Static Group tab from the IGMP page. For more information about the IGMP page, see Configuring IGMP, page 55-2.

Related Topics

- Enabling PIM and IGMP, page 55-1
- Configuring Multicast Routes, page 55-8
- Configuring PIM, page 55-11

Field Reference

Table 55-4 Static Group Tab

Element	Description
Interface	The name of the interface with which the static group is associated.
Multicast Group Address	The multicast group address to which this rule applies.

Configure IGMP Static Group Parameters Dialog Box

Use the Configure IGMP Static Group Parameters dialog box to statically assign a multicast group to an interface or to change existing static group assignments.

Navigation Path

You can access the Configure IGMP Static Group Parameters dialog box from the IGMP Page - Static Group tab. For more information, see IGMP Page - Static Group Tab, page 55-6.

Related Topics

- IGMP Page - Static Group Tab, page 55-6
- Configuring IGMP, page 55-2

Field Reference

Table 55-5 Configure IGMP Static Group Parameters Dialog Box

Element	Description
Interface	The name of the interface with which the static group is associated.
Multicast Group	The multicast group address to which this rule applies. The group address must be from 224.0.0.0 to 239.255.255.255.

IGMP Page - Join Group Tab

Use the Join Group tab to configure an interface to be a member of a multicast group.

Navigation Path

You can access the Join Group tab from the IGMP page. For more information about the IGMP page, see Configuring IGMP, page 55-2.

Related Topics

- Enabling PIM and IGMP, page 55-1
- Configuring PIM, page 55-11
- Configuring Multicast Routes, page 55-8

Field Reference

Table 55-6 Join Group Tab

Element	Description
Interface	The name of the interface for which you are configuring multicast group membership.
Multicast Group Address	The multicast group address to which this rule applies.

Configure IGMP Join Group Parameters Dialog Box

Use the Configure IGMP Join Group Parameters dialog box to configure an interface to be a member of a multicast group or to change existing membership information.

Navigation Path

You can access the Configure IGMP Join Group Parameters dialog box from the IGMP Page - Join Group tab. For more information, see IGMP Page - Join Group Tab, page 55-7.

Related Topics

- IGMP Page - Join Group Tab, page 55-7
- Configuring IGMP, page 55-2

Field Reference

Table 55-7 Configure IGMP Join Group Parameters Dialog Box

Element	Description
Interface	The name of the interface for which you are configuring multicast group membership.
Join Group	The multicast group address to which this rule applies. The group address must be from 224.0.0.0 to 239.255.255.255.

Configuring Multicast Routes

Static multicast routes let you separate multicast traffic from unicast traffic. For example, when a path between a source and destination does not support multicast routing, the solution is to configure two multicast devices with a GRE tunnel between them, sending the multicast packets over the tunnel.

Static multicast routes are local to the security appliance and are not advertised or redistributed.

Use the Multicast Routes page to manage static multicast routes—currently defined routes are listed, and you can add, edit and delete static multicast routes.

See Add/Edit MRoute Configuration Dialog Box, page 55-8 for more information about the fields displayed in the table on this page.

Navigation Path

- (Device view) Select **Platform > Multicast > Multicast Routes** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Multicast > Multicast Routes** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 55, “Configuring Multicast Policies on Firewall Devices”
- Enabling PIM and IGMP, page 55-1
- Configuring IGMP, page 55-2
- Configuring PIM, page 55-11

Add/Edit MRoute Configuration Dialog Box

Use the Add/Edit MRoute Configuration dialog box to add a static multicast route to the security appliance, or to change an existing route.

Navigation Path

You can access the Add/Edit MRoute Configuration dialog box from the Multicast Routing page. See Configuring Multicast Routes, page 55-8 for more information.

Field Reference

Table 55-8 Add/Edit MRoute Configuration Dialog Box

Element	Description
Source Interface	Enter or Select the incoming interface for the multicast route.
Source Network	Enter the IP address and mask of the multicast source, or select a Networks/Hosts object.
Output Interface/Dense	(Optional) Enter or Select the outgoing interface for the multicast route. If you specify the destination interface, the route is forwarded through the selected interface. If you do not specify a destination interface, then RPF is used to forward the route. You can specify the interface, or the RPF neighbor, but not both at the same time.
Multicast Network (PIX 6.3)	Enter or Select the group that is to receive the multicast packets. This must be a multicast IP address in the range of 224.0.1.0 to 239.255.255.255.
Distance (PIX 7.x, ASA and FWSM)	Enter an administrative distance for the static multicast route. If the static multicast route has the same administrative distance as the unicast route, the static multicast route takes precedence.

Configuring Multicast Boundary Filters

On an ASA running version 7.2(1) or later, you can use the Multicast Boundary Filter page to configure the appliance to act as a boundary between multicast domains. The ASA compares multicast group addresses to an access list, blocking all multicast traffic except that specifically permitted by the list.

The Multicast Boundary Filter page lists all currently defined per-interface boundary filter lists; you can add, edit and delete filter lists from this page.

Refer to Add/Edit MBoundary Configuration Dialog Box, page 55-9 for a description of the fields on this page.

Navigation Path

- (Device view) Select **Platform > Multicast > Multicast Boundary Filter** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Multicast > Multicast Boundary Filter** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Add/Edit MBoundary Interface Configuration Dialog Box, page 55-10

Add/Edit MBoundary Configuration Dialog Box

Use the Add/Edit MBoundary Configuration dialog box to add, edit and delete multicast boundary filter lists for individual interfaces.

Navigation Path

You can access the Add/Edit MBoundary Configuration dialog box from the Configuring Multicast Boundary Filters, page 55-9.

Related Topics

- Add/Edit MBoundary Interface Configuration Dialog Box, page 55-10
- Configuring Multicast Boundary Filters, page 55-9

Field Reference

Table 55-9 Add/Edit MBoundary Configuration Dialog Box

Element	Description
Interface	Enter or Select an interface for this multicast boundary.
Remove any Auto_RP group range announcements	If you check this box, Auto-RP messages denied by the boundary access control list for this interface are dropped. This is referred to as AutoFiltering.
Multicast boundary filter configuration list	Lists the multicast group addresses specifically permitted or denied for the specified interface. This list is managed with the Add/Edit MBoundary Interface Configuration Dialog Box, page 55-10 (click Add Row or Edit Row).

Add/Edit MBoundary Interface Configuration Dialog Box

Use this dialog box to define permit or deny multicast group entries for the list in the Add/Edit MBoundary Configuration dialog box.

Navigation Path

You can access the Add/Edit MBoundary Interface Configuration dialog box from the Add/Edit MBoundary Configuration Dialog Box, page 55-9.

Related Topics

- Configuring Multicast Boundary Filters, page 55-9

Field Reference

Table 55-10 Add/Edit MBoundary Interface Configuration Dialog Box

Element	Description
Action	Choose **permit** or **deny** to specify the action taken for this multicast group.
Multicast Group	Enter a single multicast address, or a multicast group address, to which this action applies. The address must be 0.0.0.0, or from 224.0.0.0 to 239.255.255.255. A group address range can be entered using either a standard subnet mask (e.g., 239.0.0.0 255.0.0.0), or using CIDR prefix notation (e.g., 239.0.0.0/8). You also can Select a named network/host object.

Configuring PIM

Protocol independent multicast (PIM) provides a scalable method for determining the best paths in a network for distributing a specific multicast transmission to each host that has registered using IGMP to receive the transmission. Routers and security devices use PIM to maintain tables for forwarding multicast datagrams.

With PIM sparse mode (PIM SM), which is the default for Cisco routers, when the source of a multicast transmission begins broadcasting, the traffic is forwarded from one multicast router to the next until the packets reach every registered host. If a more direct path to the traffic source exists, the last-hop router sends a join message to the source that causes the traffic to be rerouted along the better path.

Note PIM is not supported with PAT—the PIM protocol does not use ports and PAT only works with protocols that use ports.

When you enable multicast routing on a security appliance, PIM and IGMP are enabled on all interfaces by default. You can disable PIM on a per-interface basis.

The PIM page provides up to six tabbed panels:

- PIM Page - Protocol Tab, page 55-11 – Lets you manage interface-specific PIM properties.
- PIM Page - Neighbor Filter Tab, page 55-12 – Lets you manage neighbor filters for individual interfaces; available only on ASA 7.2(1)+ devices.
- PIM Page - Bidirectional Neighbor Filter Tab, page 55-13 – Lets you manage bidirectional neighbor filters for individual interfaces; available only on ASA 7.2(1)+ devices.
- PIM Page - Rendezvous Points Tab, page 55-15 – When you configure PIM, you must choose one or more devices to operate as the rendezvous point (RP). An RP is a single, common root of a shared distribution tree and is statically configured on each device. First-hop routers use the RP to send registration packets on behalf of the source multicast hosts.
- PIM Page - Route Tree Tab, page 55-17 – By default, PIM leaf routers join the shortest-path tree immediately after the first packet arrives from a new source. This reduces delay, but requires more memory than shared tree. You can configure whether the security appliance should join shortest-path tree, or use a shared tree, either for all multicast groups or only for specific multicast addresses.
- PIM Page - Request Filter Tab, page 55-18 – When the security appliance is acting as an RP, you can restrict specific multicast sources from registering. This prevents unauthorized sources from registering with the RP. The Request Filter panel lets you define the multicast sources from which the security appliance will accept PIM registration messages.

PIM Page - Protocol Tab

Use the Protocol tab to configure PIM properties for the interfaces on a security appliance (not available on PIX 6.3 devices). All currently configured interfaces are listed; you can add, edit and delete entries on this panel.

Refer to Add/Edit PIM Protocol Dialog Box, page 55-12 for a description of the fields on this panel.

Navigation Path

You access the Protocol tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Route Tree Tab, page 55-17
- PIM Page - Request Filter Tab, page 55-18

Add/Edit PIM Protocol Dialog Box

Use the Add/Edit PIM Protocol dialog box to configure PIM properties for an interface on a security appliance running PIX 7.x or later.

About the Designated Router

The DR is responsible for sending PIM register, join, and prune messaged to the Rendezvous Point (RP). When there is more than one multicast routing device on a network segment, there is an election process to select the Designated Router based on DR priority. If multiple devices have the same DR priority, then the device with the highest IP address becomes the DR. By default, security appliances have a DR priority of 1.

Navigation Path

You can access the Add/Edit PIM Protocol dialog box from the PIM Page - Protocol Tab, page 55-11.

Field Reference

Table 55-11 Add/Edit PIM Protocol Dialog Box

Element	Description
Interface	Enter or Select the interface on which you are configuring PIM.
Enable PIM	When checked, PIM is enabled on the selected interface. You can deselect this option to disable PIM on the interface without deleting this PIM Protocol entry from the table.
DR Priority	The designated router (DR) priority for this interface. The router with the highest DR priority on subnet becomes the designated router. Valid values range from 0 to 4294967294. The default DR priority is 1. Setting this value to zero makes the security appliance interface ineligible to become the default router.
Hello Interval (seconds)	The frequency, in seconds, at which the interface sends PIM hello messages. Valid values range from 1 to 3600 seconds; the default value is 30 seconds.
Join-Prune Interval (seconds)	The frequency, in seconds, at which the interface sends PIM join and prune advertisements. Valid values range from 10 to 600 seconds; the default value is 60 seconds.

PIM Page - Neighbor Filter Tab

A PIM neighbor filter is an access control list (ACL) that defines the neighbor devices that can participate in PIM. If a neighbor filter is not configured for an interface, then there are no restrictions. If a PIM neighbor filter is configured, only those neighbors permitted by the filter list can participate in PIM with the security appliance.

On an ASA running version 7.2(1) or later, you can use the Neighbor Filter tab to control the devices that can become PIM neighbors. This panel is used to define and manage the per-interface neighbor filter list. Refer to Add/Edit PIM Neighbor Filter Dialog Box, page 55-13 for a description of the fields on this panel.

Navigation Path

You access the Protocol tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Bidirectional Neighbor Filter Tab, page 55-13
- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Route Tree Tab, page 55-17
- PIM Page - Request Filter Tab, page 55-18

Add/Edit PIM Neighbor Filter Dialog Box

Use the Add/Edit PIM Neighbor Filter dialog box to add and edit entries in the PIM neighbor filter ACL displayed on the Neighbor Filter panel of the PIM page.

Navigation Path

You can access the Add/Edit PIM Neighbor Filter dialog box from the PIM Page - Neighbor Filter Tab, page 55-12.

Field Reference

Table 55-12 Add/Edit PIM Neighbor Filter Dialog Box

Element	Description
Interface	Enter or Select the interface to which this PIM Neighbor filter entry will be applied.
Neighbor Filter Group	Enter a single multicast address, or a multicast group address, to which the chosen Action applies. A group address range can be entered using either a standard subnet mask (e.g., 239.0.0.0 255.0.0.0), or using CIDR prefix notation (e.g., 239.0.0.0/8). You also can Select a named network/host object.
Action	Choose **permit** to allow the specified neighbors to participate in PIM, or **deny** to prevent the specified neighbors from participating in PIM.

PIM Page - Bidirectional Neighbor Filter Tab

A PIM bidirectional neighbor filter is an access control list (ACL) that defines the neighbor devices that can participate in the bidirectional trees and designated forwarder (DF) election. If a PIM bidirectional neighbor filter is not configured for an interface, then there are no restrictions. If a PIM bidirectional neighbor filter is configured, only those neighbors permitted by the ACL can participate in DF election process.

The PIM bidirectional neighbor filters enable the transition from a sparse-mode-only network to a "bidir" network by letting you specify the devices that should participate in DF election, while still allowing all devices to participate in the sparse-mode domain. The bidir-enabled devices can elect a DF from among themselves, even when there are non-bidir devices on the segment. Multicast boundaries on the non-bidir devices prevent PIM messages and data from the bidir groups from leaking in or out of the bidir subset cloud.

Bidirectional PIM allows multicast devices to maintain reduced state information. All of the multicast devices in a segment must be bidirectionally enabled for bidir to elect a DF.

When a PIM bidirectional neighbor filter is enabled, the routers and other devices that are permitted by the ACL are considered to be bidir-capable. Therefore:

- If a permitted neighbor does not support bidir, the DF election does not occur.
- If a denied neighbor supports bidir, then DF election does not occur.
- If a denied neighbor does not support bidir, the DF election can occur.

Managing the Bidirectional Neighbor Filter List

On an ASA running version 7.2(1) or later, you can use this panel to define and manage the per-interface bidirectional neighbor filter list, permitting or denying multicast source addresses for specific interfaces. Refer to Add/Edit PIM Bidirectional Neighbor Filter Dialog Box, page 55-14 for a description of the fields on this panel.

Navigation Path

You access the Bidirectional Neighbor Filter tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Neighbor Filter Tab, page 55-12
- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Route Tree Tab, page 55-17
- PIM Page - Request Filter Tab, page 55-18

Add/Edit PIM Bidirectional Neighbor Filter Dialog Box

Use the Add/Edit PIM Bidirectional Neighbor Filter dialog box to add or edit an entry in the bidirectional neighbor access control list displayed on the PIM Page - Bidirectional Neighbor Filter Tab, page 55-13.

Navigation Path

You can access the Add/Edit PIM Bidirectional Neighbor Filter dialog box from the PIM Page - Bidirectional Neighbor Filter Tab, page 55-13.

Field Reference

Table 55-13 Add/Edit PIM Bidirectional Neighbor Filter Dialog Box

Element	Description
Interface	Enter or Select the interface to which this PIM Bidirectional Neighbor filter entry will be applied.

Table 55-13 Add/Edit PIM Bidirectional Neighbor Filter Dialog Box (continued)

Element	Description
Neighbor Filter Group	Enter a single multicast address, or a multicast group address, to which the chosen Action applies. A group address range can be entered using either a standard subnet mask (e.g., 239.0.0.0 255.0.0.0), or using CIDR prefix notation (e.g., 239.0.0.0/8). You also can Select a named network/host object.
Action	Choose **permit** to allow the specified neighbors to participate in the DF election process, or **deny** to prevent the specified neighbors from participating in the process.

PIM Page - Rendezvous Points Tab

When you configure PIM, you must choose one or more routers or routing devices to operate as the RP. An RP is a single, common root of a shared distribution tree and is statically configured on each device. First hop routers use the RP to send register packets on behalf of the source multicast hosts.

You can configure a single RP to serve more than one group. If a specific group is not specified, the RP for the group is applied to the entire IP multicast group range (224.0.0.0/4).

Use the Rendezvous Points panel to define rendezvous points. You can configure more than one RP, but you cannot have more than one entry with the same RP.

Navigation Path

You access the Rendezvous Points tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Route Tree Tab, page 55-17
- PIM Page - Request Filter Tab, page 55-18

Field Reference

Table 55-14 Rendezvous Points Tab

Element	Description
Generate older IOS compatible register messages	Check this box if your rendezvous point is a Cisco IOS router. The security appliance software accepts register messages with the checksum calculated on the PIM header and only the next 4 bytes, while Cisco IOS software accepts register messages with the checksum calculated on the entire PIM message for all PIM message types.
Rendezvous Points table	Lists the rendezvous points currently configured on the security appliance. Use the Add Row, Edit Row and Delete Row buttons to manage this list; the Add Row and Edit Row buttons open the Add/Edit Rendezvous Point Dialog Box, page 55-16.

Add/Edit Rendezvous Point Dialog Box

Use the Add/Edit Rendezvous Point dialog box to add an entry to the Rendezvous Points table, or to edit an existing rendezvous point entry. Please note the following:

- You cannot use the same rendezvous point address twice.
- You cannot specify “All Groups” for more than one rendezvous point.

Navigation Path

You can access the Add/Edit Rendezvous Point dialog box from the PIM Page - Rendezvous Points Tab, page 55-15.

Field Reference

Table 55-15 Add/Edit Rendezvous Point Dialog Box

Element	Description
Rendezvous Point IP Address	Enter the IP address of the rendezvous point. This is a unicast address. You also can click Select to select a Networks/Hosts object. When editing a rendezvous point entry, you cannot change this value.
Use bi-directional forwarding	Check this box if you want the specified Multicast Groups to operate in bidirectional mode. In bidirectional mode, if the security appliance receives a multicast packet and has no directly connected members or PIM neighbors present, it sends a Prune message back to the source. Deselect this option if you want the specified Multicast Groups to operate in Sparse Mode. **Note** The security appliance always advertises bidirectional capability in PIM hello messages regardless of the actual bidir configuration.
Use this RP for All Multicast Groups	Select this option to use the specified Rendezvous Point for all multicast groups on the interface.
Use this RP for the Multicast Groups as specified below	Select this option to define the multicast groups that are to use the specified Rendezvous Point; the Multicast Groups table is activated.
Multicast Groups table	The multicast groups currently associated with the specified Rendezvous Point are listed. Table entries are processed from the top down. For example, you can create an entry that includes a range of multicast groups, and then exclude specific groups within that range by placing deny rules for those specific groups at the top of the table. That is, the permit rule for the range of multicast groups follows the individual deny statements. Use the buttons at the bottom of the table to open the Add/Edit Multicast Group Rules Dialog Box, page 55-17 to add or edit an entry; to delete an entry; and to move entries up or down in the table.

Add/Edit Multicast Group Rules Dialog Box

Use the Add/Edit Multicast Group Rules dialog box to create a multicast group rule, or modify a multicast group rule, for the Multicast Groups table in the Add/Edit Rendezvous Point dialog box. This dialog box is also used to specify individual multicast groups that use Shared Tree route filtering on the Route Tree tab

Navigation Path

When defining Rendezvous Points, you access the Add/Edit Multicast Group Rules dialog box from the Add/Edit Rendezvous Point Dialog Box, page 55-16. See PIM Page - Rendezvous Points Tab, page 55-15 for more information.

When specifying how PIM register messages are filtered, you open this dialog box by clicking Add Row or Edit row buttons below the Multicast Groups table on the PIM Page - Route Tree Tab, page 55-17.

Field Reference

Table 55-16 Add/Edit Multicast Group Rules Dialog Box

Element	Description
Action	Choose **permit** to create a group rule that allows the specified multicast addresses; choose **deny** to create a group rule that denies the specified multicast addresses.
Multicast Group Network	Enter the multicast address and network mask associated with the group, or Select the desired Networks/Hosts object.

PIM Page - Route Tree Tab

If the security appliance is acting as a Rendezvous Point, use the Route Tree tab to specify how the PIM register messages from various sources are filtered: shortest-path tree or shared tree, either for all multicast groups or only for specific multicast addresses.

Navigation Path

You can access the Route Tree tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Request Filter Tab, page 55-18

Field Reference

Table 55-17 Route Tree Tab

Element	Description
If.., specify how the PIM register messages from various sources are filtered	Select a tree/groups option: • **Use Shortest Path Tree for All Groups** – The security appliance uses shortest-path tree for all multicast groups. • **Use Shared Tree for All Groups** – The security appliance uses shared tree for all multicast groups. • **Use Shared Tree for the Groups specified below** – The security appliance uses shared tree for those groups specified below in the Multicast Groups table. Shortest-path tree is used for any group not listed in the Multicast Groups table.
Multicast Groups table	The multicast groups using Shared Tree are listed. Table entries are processed from the top down. For example, you can create an entry that includes a range of multicast groups, and then exclude specific groups within that range by placing deny rules for those specific groups at the top of the table. That is, the permit rule for the range of multicast groups follows the individual deny statements. Use the buttons at the bottom of the table to open the Add/Edit Multicast Group Rules Dialog Box, page 55-17 to add or edit an entry; to delete an entry; and to move entries up or down in the table.

PIM Page - Request Filter Tab

When the security appliance acts as a rendezvous point, you can restrict specific multicast sources from registering with it. This prevents unauthorized sources from registering with the rendezvous point. You can use the Request Filter tab to define the multicast sources from which the security appliance accepts and denies PIM register messages.

Navigation Path

You can access the Request Filter tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Route Tree Tab, page 55-17

Field Reference

Table 55-18 Request Filter Tab

Element	Description
Filter PIM register messages using	Choose how PIM register messages are filtered for different multicast groups: • **None** – Do not filter PIM register messages. • **route-map** – Filter PIM register messages using a specified route map; the Route Map field is activated. Only PIM register messages that are permitted by the route map are allowed to reach the rendezvous point. • **access-list** – Filter PIM register messages using an access list; the Multicast Groups table is activated. Only PIM register messages that are permitted by the access list are allowed to reach the rendezvous point.
Route Map	When **route-map** is the chosen filter, enter a route-map name. Use standard host ACLs in the referenced route map; extended ACLs are not supported. **Note** This field contains only the Route Map name. The Route Map is created and contained within a FlexConfig; see Chapter 7, "Managing FlexConfigs" for more information.
Multicast Groups table	Lists the currently defined multicast group Request Filter rules. Table entries are processed from the top down. For example, you can create an entry that includes a range of multicast groups, and then exclude specific groups within that range by placing deny rules for those specific groups at the top of the table. That is, the permit rule for the range of multicast groups follows the individual deny statements. Use the buttons at the bottom of the table to open the Add/Edit Multicast Group Rules Dialog Box, page 55-19 to add or edit an entry; to delete an entry; and to move entries up or down in the table.

Add/Edit Multicast Group Rules Dialog Box

Use the Add/Edit Multicast Group Rules dialog box to define the multicast sources that are denied or permitted to register with the security appliance when the appliance acts as a rendezvous point. You create the filter rules based on the source IP address and the destination multicast address.

Navigation Path

You can access the Add/Edit Multicast Group Rules dialog box from the PIM Page - Request Filter Tab, page 55-18.

Field Reference

Table 55-19 Add/Edit Multicast Group Rules Dialog Box

Element	Description
Action	Choose **permit** to create a rule that allows the specified Source of the specified Destination multicast traffic to register with the security appliance; choose **deny** to create a rule that denies registration to the specified Source/Destination multicast traffic.
Source Network	Enter the IP address and network mask for the source of the register message, or Select the appropriate Networks/Hosts object.
Destination Network	Enter the IP address and network mask for the multicast destination, or Select the appropriate Networks/Hosts object.

PIM Page - Bootstrap Router Tab

PIM Bootstrap Router (BSR) is a dynamic Rendezvous Point (RP) selection model that uses candidate routers for Rendezvous Point function and for relaying the Rendezvous Point information for a group.The Rendezvous Point function includes RP discovery and provides a default route to the RP. It does this by configuring a set of devices as candidate BSRs (C-BSR) which participate in a BSR election process to choose a BSR amongst themselves. Once the BSR is chosen, devices that are configured as candidate Rendezvous Points (C-RP) will start sending their group mapping to the elected BSR. The BSR then distributes the group-to-RP mapping information to all the other devices down the multicast tree, through BSR messages that travel from PIM router to PIM router on a per-hop basis.

You can use the Boot Strap Router tab to configure a device as a PIM Bootstrap Router.

Navigation Path

You can access the Bootstrap Router tab from the PIM page. For more information, see Configuring PIM, page 55-11.

Related Topics

- PIM Page - Protocol Tab, page 55-11
- PIM Page - Rendezvous Points Tab, page 55-15
- PIM Page - Route Tree Tab, page 55-17

Field Reference

Table 55-20 Bootstrap Router Tab

Element	Description
Interface	Select the interface from which the BSR address is derived, to make it a candidate BSR.
Hash Mask Length	Enter the length of a mask that is to be ANDed with the group address before the hash function is called. All groups with the same seed hash (correspond) to the same Rendezvous Point (RP). For example, if this value is 24, only the first 24 bits of the group addresses matter. This fact allows you to get one RP for multiple groups.

Table 55-20 Bootstrap Router Tab (continued)

Element	Description
Priority	Enter the priority of the candidate BSR. The BSR with the larger priority is preferred. If the priority values are the same, the router with the larger IP address is the BSR.

Add/Edit Bootstrap Router Dialog Box

Use the Add/Edit Bootstrap Router dialog box to configure a Bootstrap Router as a Border BSR. BSR messages should not be exchanged between different domains, because routers in one domain may elect rendezvous points (RPs) in the other domain, resulting in protocol malfunction or loss of isolation between the domains

A border interface in a PIM sparse mode (PIM-SM) domain is configured to avoid exchange of certain traffic with a neighboring domain reachable through that interface, especially if that domain is also running PIM-SM. Thus to prevent BSR messages from being sent or received through such an interface configure the interface as a border BSR.

Navigation Path

You can access the Add/Edit Bootstrap Router dialog box from the PIM Page - Bootstrap Router Tab, page 55-20

Field Reference

Table 55-21 Add/Edit Bootstrap Router Dialog Box

Element	Description
(Optional) BSR Border table Add	Add an interface and configure it as a border BSR. No PIM BSR messages will be sent or received, when an interface is configured as a Border BSR.

CHAPTER 56

Configuring Routing Policies on Firewall Devices

The Routing section in Security Manager contains pages for defining and managing routing settings for security appliances.

This chapter contains the following topics:

- Configuring No Proxy ARP, page 56-1
- Configuring BGP, page 56-2
- Configuring EIGRP, page 56-32
- Configuring ISIS, page 56-49
- Configuring BFD Routing, page 56-66
- Configuring OSPF, page 56-75
- Configuring OSPFv3, page 56-100
- Configuring RIP, page 56-119
- Configuring Static Routes, page 56-128
- Configuring Policy Objects for ASA Routing Policies, page 56-132

Configuring No Proxy ARP

When a host sends IP traffic to another device on the same Ethernet network, the host needs to know the MAC address of the device. Address Resolution Protocol (ARP) is a Layer 2 protocol that resolves an IP address to a MAC address: a host sends an ARP request asking “Who is this IP address?” The device owning the IP address replies, “I own that IP address; here is my MAC address.”

With Proxy ARP, a device responds to an ARP request with its own MAC address, even though the device does not own the IP address. Serving as an ARP Proxy for another host effectively directs network traffic to the proxy, in this case your security appliance. Traffic that passes through the appliance is then routed to the appropriate destination.

For example, the security appliance uses proxy ARP when you configure NAT and specify a global address that is on the same network as the appliance interface. The only way traffic can reach the destination hosts is if the appliance claims and subsequently routes traffic to the destination global addresses.

By default, proxy ARP is enabled for all interfaces. Use the No Proxy ARP page to disable proxy ARP for global addresses:

- To disable proxy ARP for one or more interfaces, enter their names in the Interfaces field. Separate multiple interfaces with commas. You can click Select to choose the interfaces from a list of interfaces defined on the device, and interface roles defined in Security Manager.

Note On ASA 8.4.2 and later devices operating in routed mode, you can disable Proxy ARP on the egress interface for a Manual NAT rule. See Do not proxy ARP on Destination Interface for more information.

Navigation Path

- (Device view) Select **Platform > Routing > No Proxy ARP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > No Proxy ARP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Configuring Static Routes, page 56-128
- Configuring RIP, page 56-119
- Configuring OSPF, page 56-75

Configuring BGP

Border Gateway Protocol (BGP) is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

Note BGP configuration is supported on ASA 9.2(1)+ only. Also, beginning with ASA 9.3(1), BGP is supported in L2 (EtherChannel Type) and L3 (Individual Interface Type) clustering modes.

Navigation Path

- (Device view) Select **Platform > Routing > BGP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > BGP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

The BGP page provides two tabbed panels for configuring BGP routing on a firewall device. This is the basic procedure for configuring the BGP process:

1. Enable the BGP routing process by checking the Enable BGP check box on the BGP page.
2. In the AS Number field, enter the autonomous system (AS) number for the BGP process. The AS number internally includes multiple autonomous numbers. The AS number can be from 1 to 4294967295 or from 1.0 to 65535.65535.
3. On the General Tab, page 56-5:
   - (Optional) Check the Limit the number of AS numbers in AS_PATH attribute of received routes check box to restrict the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
   - (Optional) Check the Log Neighbor Changes check box to enable logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.

- (Optional) Check the Use TCP path MTU Discovery check box to use the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
- (Optional) Check the Enable fast external failover check box to reset the external BGP session immediately upon link failure.
- (Optional) Check the Enforce that the first AS is peer's AS for EBGP routes check box to discard incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
- (Optional) Check the Use dot notation for AS numbers check box to split the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65553 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.
- Define the configuration related to the best path selection process for BGP routing (see General Tab, page 56-5).
- Specify the timer information in the Neighbor timers area (see General Tab, page 56-5).
- (Optional) Configure Graceful Restart (see General Tab, page 56-5).

4. On the IPv4 Family tab, select the Enable IPv4 Family check box and then use the tabs provided to configure IPv4 Address Family settings. For more information, see IPv4 Family Tab, page 56-6.
5. On the IPv6 Family tab, select the Enable IPv6 Family check box and then use the tabs provided to configure IPv6 Address Family settings. For more information, see IPv6 Family Tab, page 56-20

Related Topics

- About BGP, page 56-3

About BGP

BGP is an inter autonomous system routing protocol. An autonomous system is a network or group of networks under a common administration and with common routing policies. BGP is used to exchange routing information for the Internet and is the protocol used between Internet service providers (ISP).

When to Use BGP

Customer networks, such as universities and corporations, usually employ an Interior Gateway Protocol (IGP) such as OSPF for the exchange of routing information within their networks. Customers connect to ISPs, and ISPs use BGP to exchange customer and ISP routes. When BGP is used between autonomous systems (AS), the protocol is referred to as External BGP (EBGP). If a service provider is using BGP to exchange routes within an AS, then the protocol is referred to as Interior BGP (IBGP).

Routing Table Changes

BGP neighbors exchange full routing information when the TCP connection between neighbors is first established. When changes to the routing table are detected, the BGP routers send to their neighbors only those routes that have changed. BGP routers do not send periodic routing updates, and BGP routing updates advertise only the optimal path to a destination network.

Routes learned via BGP have properties that are used to determine the best route to a destination, when multiple paths exist to a particular destination. These properties are referred to as BGP attributes and are used in the route selection process:

- Weight -- This is a Cisco-defined attribute that is local to a router. The weight attribute is not advertised to neighboring routers. If the router learns about more than one route to the same destination, the route with the highest weight is preferred.
- Local preference -- The local preference attribute is used to select an exit point from the local AS. Unlike the weight attribute, the local preference attribute is propagated throughout the local AS. If there are multiple exit points from the AS, the exit point with the highest local preference attribute is used as an exit point for a specific route.
- Multi-exit discriminator -- The multi-exit discriminator (MED) or metric attribute is used as a suggestion to an external AS regarding the preferred route into the AS that is advertising the metric. It is referred to as a suggestion because the external AS that is receiving the MEDs may also be using other BGP attributes for route selection. The route with the lower MED metric is preferred.
- Origin -- The origin attribute indicates how BGP learned about a particular route. The origin attribute can have one of three possible values and is used in route selection.
  - IGP- The route is interior to the originating AS. This value is set when the network router configuration command is used to inject the route into BGP.
  - EGP-The route is learned via the Exterior Border Gateway Protocol (EBGP).
  - Incomplete- The origin of the route is unknown or learned in some other way. An origin of incomplete occurs when a route is redistributed into BGP.
- AS_path -- When a route advertisement passes through an autonomous system, the AS number is added to an ordered list of AS numbers that the route advertisement has traversed. Only the route with the shortest AS_path list is installed in the IP routing table.
- Next hop -- The EBGP next-hop attribute is the IP address that is used to reach the advertising router. For EBGP peers, the next-hop address is the IP address of the connection between the peers. For IBGP, the EBGP next-hop address is carried into the local AS.
- Community -- The community attribute provides a way of grouping destinations, called communities, to which routing decisions (such as acceptance, preference, and redistribution) can be applied. Route maps are used to set the community attribute. The predefined community attributes are as follows:
  - no-export- Do not advertise this route to EBGP peers.
  - no-advertise- Do not advertise this route to any peer.
  - internet- Advertise this route to the Internet community; all routers in the network belong to it.

BGP Path Selection

BGP may receive multiple advertisements for the same route from different sources. BGP selects only one path as the best path. When this path is selected, BGP puts the selected path in the IP routing table and propagates the path to its neighbors. BGP uses the following criteria, in the order presented, to select a path for a destination:

- If the path specifies a next hop that is inaccessible, drop the update.
- Prefer the path with the largest weight.
- If the weights are the same, prefer the path with the largest local preference.
- If the local preferences are the same, prefer the path that was originated by BGP running on this router.
- If no route was originated, prefer the route that has the shortest AS_path.
- If all paths have the same AS_path length, prefer the path with the lowest origin type (where IGP is lower than EGP, and EGP is lower than incomplete).

- If the origin codes are the same, prefer the path with the lowest MED attribute.
- If the paths have the same MED, prefer the external path over the internal path.
- If the paths are still the same, prefer the path through the closest IGP neighbor.
- If both paths are external, prefer the path that was received first (the oldest one).
- Prefer the path with the lowest IP address, as specified by the BGP router ID.
- If the originator or router ID is the same for multiple paths, prefer the path with the minimum cluster list length.
- Prefer the path that comes from the lowest neighbor address.

General Tab

Use the General tab to configure BGP settings such as Best Path Selection, Neighbor Timers, and Graceful Restart.

Navigation Path

You can access the Neighbors tab from the BGP page (see Configuring BGP, page 56-2).

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family Tab, page 56-6

Field Reference

Table 56-1 General Tab

Element	Description
Limit the number of AS numbers in AS_PATH attribute of received routes	Restricts the number of AS numbers in AS_PATH attribute to a specific number. Valid values are from 1 to 254.
Log Neighbor Changes	Enables logging of BGP neighbor changes (up or down) and resets. This helps in troubleshooting network connectivity problems and measuring network stability.
Use TCP path MTU Discovery	Enables the use of the Path MTU Discovery technique to determine the maximum transmission unit (MTU) size on the network path between two IP hosts. This avoids IP fragmentation.
Enable fast external failover	Resets the external BGP session immediately upon link failure.
Enforce that the first AS is peer's AS for EBGP routes	Discards incoming updates received from external BGP peers that do not list their AS number as the first segment in the AS_PATH attribute. This prevents a mis-configured or unauthorized peer from misdirecting traffic by advertising a route as if it was sourced from another autonomous system.
Use dot notation for AS numbers	Splits the full binary 4-byte AS number into two words of 16 bits each, separated by a dot. AS numbers from 0-65553 are represented as decimal numbers and AS numbers larger than 65535 are represented using the dot notation.

Table 56-1 General Tab (continued)

Element	Description
Best Path Selection	
Default local preference	Specify a value between 0 and 4294967295. The default value is 100. Higher values indicate higher preference. This preference is sent to all routers and access servers in the local autonomous system.
Allow comparing MED from different neighbors	Allows the comparison of Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems.
Compare Router-id for identical EBGP paths	Compares similar paths received from external BGP peers during the best path selection process and switches the best path to the route with the lowest router ID.
Pick the best MED path among paths advertised from the neighboring AS	Enables MED comparison among paths learned from confederation peers. The comparison between MEDs is made only if no external autonomous systems are there in the path.
Treat missing MED as the least preferred one	Considers the missing MED attribute as having a value of infinity, making the path the least desirable; therefore, a path with a missing MED is least preferred.
Neighbor Timers	
Keepalive Interval	Enter the time interval for which the BGP neighbor remains active after not sending a keepalive message. At the end of this keepalive interval, the BGP peer is declared dead, if no messages are sent. The default value is 60 seconds.
Hold Time	Enter the time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. The default values is 180 seconds.
Min Hold Time	(Optional) Enter the minimum time interval for which the BGP neighbor remains active while a BGP connection is being initiated and configured. Specify a value from 0 to 65535.
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	
Enable Graceful Restart	Enables ASA peers to avoid a routing flap following a switchover.
Restart Time	Specify the time duration that ASA peers will wait to delete stale routes before a BGP open message is received. The default value is 120 seconds. Valid values are between 1 and 3600 seconds.
Stalepath Time	Enter the time duration that the ASA will wait before deleting stale routes after an end of record (EOR) message is received from the restarting ASA. The default value is 360 seconds. Valid values are between 1 and 3600 seconds.

IPv4 Family Tab

Use the IPv4 Family tab on the BGP page to enable and configure IPv4 settings for BGP.

Navigation Path

You can access the IPv4 Family tab from the BGP page. For more information about the BGP page, see Configuring BGP, page 56-2.

Related Topics

- About BGP, page 56-3
- General Tab, page 56-5

Field Reference

Table 56-2 IPv4 Family - Aggregate Address Tab

Element	Description
Enable IPv4 Family	Enables configuration of routing sessions that use standard IPv4 address prefixes.
General	Use this panel to configure general IPv4 settings such as Best Path Selection, Neighbor Timers, and Graceful Restart. See IPv4 Family - General Tab, page 56-7 for more about these definitions.
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box, page 56-9 for more about these definitions.
Filtering	Use this panel to filter routes or networks received in incoming BGP updates. See Add/Edit Filter Dialog Box, page 56-10 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box, page 56-11 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box, page 56-17 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box, page 56-18 for more about these definitions.
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box, page 56-19 for more about these definitions.

IPv4 Family - General Tab

Use the IPv4 Family - General tab to configure the general IPv4 settings.

Navigation Path

You can access the General tab from the IPv4 Family Tab on the BGP page. For more information about the IPv4 Family tab, see IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3

Field Reference

Table 56-3 IPv4 Family - General Tab

Element	Description
Router ID	On a single device, choose **Automatic** or **IP Address**. (An address field appears when you choose IP Address.) If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field. On a device cluster, choose **Automatic** or **Cluster Pool**. (An IPv4 Pool object ID field appears when you choose Cluster Pool.) If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see Add or Edit IPv4 Pool Dialog Box, page 6-92.
Learned Route Map	Enter or Select the name of a route map object. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Scanning Interval	Enter a scanning interval (in seconds) for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.
Routes and Synchronization	
Generate Default Route	(Optional) Configures a BGP routing process to distribute a default route (network 0.0.0.0).
Summarize subnet routes into network-level routes	(Optional) Configures automatic summarization of subnet routes into network-level routes.
Advertise inactive routes	(Optional) Advertises routes that are not installed in the routing information base (RIB).
Synchronize between BGP and the Interior Gateway Protocol (IGP) system	Enables synchronization between BGP and your Interior Gateway Protocol (IGP) system. To enable the Cisco IOS software to advertise a network route without waiting for the IGP, deselect this option. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. By default, synchronization between BGP and the IGP is turned off to allow the Cisco IOS software to advertise a network route without waiting for route validation from the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems. Use synchronization if routers in the autonomous system do not speak BGP.

Table 56-3 IPv4 Family - General Tab (continued)

Element	Description
Redistribute iBGP into an IGP	(Optional) Configures iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
Administrative Route Distances	
External	Specifies the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
Internal	Specifies administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
Local	Specifies administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.
Next Hop	
Enable address tracking	(Optional) Enables BGP next hop address tracking.
Delay Interval	Specify the delay interval between checks on updated next-hop routes installed in the routing table.
Forward packets over Multiple Paths	
Number of Paths	(Optional) Specify the maximum number of external BGP routes that can be installed to the routing table.
IBGP Number of Paths	(Optional) Specify the maximum number of internal BGP routes that can be installed to the routing table.

Add/Edit Aggregate Address Dialog Box

Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Navigation Path

You can access the Add/Edit Aggregate Address dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Filter Dialog Box, page 56-10
- Add/Edit Neighbor Dialog Box, page 56-11
- Add/Edit Network Dialog Box, page 56-17
- Add/Edit Redistribution Dialog Box, page 56-18

- Add/Edit Route Injection Dialog Box, page 56-19

Field Reference

Table 56-4 Add/Edit Aggregate Address Dialog Box

Element	Description
Network	Enter an IP address, or enter or Select the desired Network/Hosts objects.
Attribute Map	(Optional) Enter or Select the route map used to set the attribute of the aggregate route. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Advertise Map	(Optional) Enter or Select the route map used to select the routes to create AS_SET origin communities. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Suppress Map	(Optional) Enter or Select the route map used to select the routes to be suppressed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Generate AS Set Path Information	Enables generation of autonomous system set path information.
Filter all more-specific routes from updates	Filters all more-specific routes from updates.

Add/Edit Filter Dialog Box

Use the Add/Edit Filter dialog box to filter routes or networks received in incoming BGP updates.

Navigation Path

You can access the Add/Edit Filter dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Aggregate Address Dialog Box, page 56-9

- Add/Edit Neighbor Dialog Box, page 56-11
- Add/Edit Network Dialog Box, page 56-17
- Add/Edit Redistribution Dialog Box, page 56-18
- Add/Edit Route Injection Dialog Box, page 56-19

Field Reference

Table 56-5 Add/Edit Filter Dialog Box

Element	Description
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.
Direction	Choose a direction from the Direction drop-down list. The direction will specify if the filter should be applied to inbound updates or outbound updates.
Protocol	Select the routing process for which you want to filter: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
AS Number	Shows the autonomous system number of the BGP routing process. This value is specified on the BGP page (see Configuring BGP, page 56-2).
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.

Add/Edit Neighbor Dialog Box

Use the Add/Edit Neighbor dialog box to define BGP neighbors and neighbor settings.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Aggregate Address Dialog Box, page 56-9
- Add/Edit Filter Dialog Box, page 56-10
- Add/Edit Network Dialog Box, page 56-17
- Add/Edit Redistribution Dialog Box, page 56-18
- Add/Edit Route Injection Dialog Box, page 56-19

Field Reference

Table 56-6 Add/Edit Neighbor Dialog Box

Element	Description
General	
IP Address	Enter the BGP neighbor IP address. This IP address is added to the BGP neighbor table.
Remote AS	Enter the autonomous system to which the BGP neighbor belongs.
Enable Address Family	(Optional) Enables communication with the BGP neighbor.
Shutdown neighbor administratively	(Optional) Disable a neighbor or peer group.
Configure Graceful Restart per neighbor (ASA 9.3.1+ only)	(Optional) Enables configuration of the Border Gateway Protocol (BGP) graceful restart capability for this neighbor. After selecting this option, you must use the Graceful Restart (Use in failover or spanned cluster mode) option to specify whether graceful restart should be enabled or disabled for this neighbor.
Graceful Restart (Use in failover or spanned cluster mode) (ASA 9.3.1+ only)	(Optional) Enables the Border Gateway Protocol (BGP) graceful restart capability for this neighbor.
Description	(Optional) Enter a description for the BGP neighbor.
fall-over BFD	(Optional) Enables BFD support for fall-over for the BGP neighbor.
BFD-Hop	(Optional) Specify if there is a single IP hop or multiple IP hops between a BFD source and destination.
Filtering	
Filter routes using an access list	(Optional) Enter or Select the appropriate incoming or outgoing access control list to distribute BGP neighbor information.
Filter routes using route map	(Optional) Enter or Select the appropriate incoming or outgoing route maps to apply a route map to incoming or outgoing routes. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Filter routes using a Prefix list	(Optional) Enter or Select the appropriate incoming or outgoing prefix list to distribute BGP neighbor information. **Tip** Click **Select** to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box, page 56-146.

Table 56-6 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Filter routes using AS Path filter	(Optional) Enter or Select the appropriate incoming or outgoing AS path filter to distribute BGP neighbor information. **Tip** Click **Select** to open the AS Path Object Selector from which you can select an AS path object. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes, page 56-151.
Limit the number of prefixes allowed from the neighbor	(Optional) Select to control the number of prefixes that can be received from a neighbor. • Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum Prefixes** field. • Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold Level** field. Valid values are integers between 1 and 100. The default value is 75. • (Optional) Check the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following: – Select **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field. – Select **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
Routes	
Advertisement Interval	Enter the minimum interval (in seconds) between the sending of BGP routing updates. Valid values are between 1 and 600.
Remove private AS numbers from outbound routing updates	(Optional) Excludes the private AS numbers from being advertised on outbound routes.
Generate Default route	(Optional) Select to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the Route map field. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.

Table 56-6 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Conditionally Advertised Routes	(Optional) To add or edit conditionally advertised routes, click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button. In the Add/Edit Advertised Route dialog box, do the following: • Click **Select** to open the Route Map Object Selector from which you can select a route map that will be advertised if the conditions of the exist map or the non-exist map are met. For more information about route maps, see Understanding Route Map Objects, page 56-132 • Do one of the following: – Select Set Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised. – Select Non-Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
Timers	
Set timers for the BGP peer	(Optional) Select to set the keepalive frequency, hold time and minimum hold time.
Keepalive Interval	Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
Hold Time	Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
Min Hold Time	(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 0 seconds.

Table 56-6 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Advanced	
Enable Authentication	(Optional) Select to enable MD5 authentication on a TCP connection between two BGP peers. • Choose an encryption type from the Enable Encryption drop-down list. • Enter a password in the Password field. Reenter the password in the Confirm field. The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces. **Note** You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.
Send Community attribute to this neighbor	(Optional) Specifies that communities attributes should be sent to the BGP neighbor.
Use ASA as next hop for neighbor	(Optional) Select to configure the router as the next-hop for a BGP speaking neighbor or peer group.
Disable connection verification	(Optional) Select to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IP address. This command is required only when the neighbor ebgp-multihop command is configured with a TTL value of 1. The address of the single-hop eBGP peer must be reachable. The neighbor update-source command must be configured to allow the BGP routing process to use the loopback interface for the peering session. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.
Allow connections with neighbor that is not directly connected	Select to accept and attempt BGP connections to external peers residing on networks that are not directly connected. (Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255. **Note** This feature should be used only under the guidance of Cisco technical support staff. To prevent the creation of loops through oscillating routes, the multihop will not be established if the only route to the multihop peer is the default route (0.0.0.0).

Table 56-6 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Limit number of TTL hops to neighbor	Select this option to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254. This feature provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers. This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network. This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary. To maximize the effectiveness of this feature, the hop-count value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session. The following restrictions apply to the configuration of this command: • This feature is not supported for internal BGP (iBGP) peers. • The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack. • This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.
Use TCP Path MTU Discovery	(Optional) Select to enable a TCP transport session for a BGP session.

Table 56-6 Add/Edit Neighbor Dialog Box (continued)

Element	Description
TCP transport mode	Choose the TCP connection mode from the drop-down list. Options are Default, Active, or Passive.
Weight	(Optional) Enter a weight for the BGP neighbor connection.
BGP Version	Choose the BGP version that the ASA will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.
Migration **Note** This customization should only be used for AS migration, and should be removed after the transition has been completed. The procedure should be attempted only by an experienced network operator. Routing loops can be created through improper configuration.	
Customize the AS number for routes received from the neighbor	(Optional) Select to customize the AS_PATH attribute for routes received from an eBGP neighbor.
Local AS Number	Enter the local autonomous system number. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
Do not prepend local AS number to routes received from neighbor	(Optional) Select to prevent the local AS number from being prepended to any routes received from eBGP peer.
Replace real AS number with local AS number in routes received from neighbor	(Optional) Select to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
Accept either real AS number or local AS number in routes received from neighbor	(Optional) Configures the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Add/Edit Network Dialog Box

Use the Add/Edit Network dialog box to define the networks to be advertised by the BGP routing process.

Navigation Path

You can access the Add/Edit Network dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Aggregate Address Dialog Box, page 56-9
- Add/Edit Filter Dialog Box, page 56-10
- Add/Edit Neighbor Dialog Box, page 56-11

- Add/Edit Redistribution Dialog Box, page 56-18
- Add/Edit Route Injection Dialog Box, page 56-19

Field Reference

Table 56-7 Add/Edit Network Dialog Box

Element	Description
Network	Specifies the network to be advertised by the BGP routing processes.
Route Map	(Optional) Enter or Select a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.

Add/Edit Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to define the conditions for redistributing routes from another routing domain into BGP.

Navigation Path

You can access the Add/Edit Redistribution dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Aggregate Address Dialog Box, page 56-9
- Add/Edit Filter Dialog Box, page 56-10
- Add/Edit Neighbor Dialog Box, page 56-11
- Add/Edit Network Dialog Box, page 56-17
- Add/Edit Route Injection Dialog Box, page 56-19

Field Reference

Table 56-8 Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.
Metric	(Optional) Enter a metric for the redistributed route.

Table 56-8 Add/Edit Redistribution Dialog Box (continued)

Element	Description
Route Map	Enter or Select a route map that should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Match	The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol. • Internal • External 1 • External 2 • NSSA External 1 • NSSA External 2

Add/Edit Route Injection Dialog Box

Use the Add/Edit Route Injection dialog box to define the routes to be conditionally injected into the BGP routing table.

Navigation Path

You can access the Add/Edit Route Injection dialog box from the IPv4 Family Tab, page 56-6.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv4 Family - General Tab, page 56-7
- Add/Edit Aggregate Address Dialog Box, page 56-9
- Add/Edit Filter Dialog Box, page 56-10
- Add/Edit Neighbor Dialog Box, page 56-11
- Add/Edit Network Dialog Box, page 56-17
- Add/Edit Redistribution Dialog Box, page 56-18

Field Reference

Table 56-9 Add/Edit Route Injection Dialog Box

Element	Description
Inject Map	Enter or Select the route map that specifies the prefixes to inject into the local BGP routing table. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Exist Map	Enter or Select the route map containing the prefixes that the BGP speaker will track. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Injected routes will inherit the attributes of the aggregate route	Configures the injected route to inherit attributes of the aggregate route.

IPv6 Family Tab

Use the IPv6 Family tab on the BGP page to enable and configure IPv6 settings for BGP.

Navigation Path

You can access the IPv6 Family tab from the BGP page. For more information about the BGP page, see Configuring BGP, page 56-2.

Related Topics

- About BGP, page 56-3
- General Tab, page 56-5

Field Reference

Table 56-10 IPv6 Family - Aggregate Address Tab

Element	Description
Enable IPv6 Family	Enables configuration of routing sessions that use standard IPv6 address prefixes.
General	Use this panel to configure general IPv6 settings. See IPv6 Family - General Tab, page 56-21 for more about these definitions.

Table 56-10 IPv6 Family - Aggregate Address Tab (continued)

Element	Description
Aggregate Address	Use this panel to define the aggregation of specific routes into one route. Specify a value for the aggregate timer (in seconds) in the Aggregate Timer field. Valid values are 0 or any value between 6 and 60. The default value is 30. See Add/Edit Aggregate Address Dialog Box, page 56-22 for more about these definitions.
Neighbor	Use this panel to define BGP neighbors and neighbor settings. See Add/Edit Neighbor Dialog Box, page 56-24 for more about these definitions.
Networks	Use this panel to define the networks to be advertised by the BGP routing process. See Add/Edit Network Dialog Box, page 56-29 for more about these definitions.
Redistribution	Use this panel to define the conditions for redistributing routes from another routing domain into BGP. See Add/Edit Redistribution Dialog Box, page 56-30 for more about these definitions.
Route Injection	Use this panel to define the routes to be conditionally injected into the BGP routing table. See Add/Edit Route Injection Dialog Box, page 56-31 for more about these definitions.

IPv6 Family - General Tab

Use the IPv6 Family - General tab to configure the general IPv6 settings.

Navigation Path

You can access the General tab from the IPv6 Family Tab on the BGP page. For more information about the IPv6 Family tab, see IPv6 Family Tab, page 56-20.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3

Field Reference

Table 56-11 IPv6 Family - General Tab

Element	Description
Scanning Interval	Enter a scanning interval (in seconds) for BGP routers for next-hop validation. Valid values are from 5 to 60 seconds. The default value is 60.
Routes and Synchronization	
Generate Default Routes	(Optional) Configures a BGP routing process to distribute a default route (network 0.0.0.0).

Table 56-11 IPv6 Family - General Tab (continued)

Element	Description
Advertise inactive routes	(Optional) Advertises routes that are not installed in the routing information base (RIB).
Synchronize between BGP and the Interior Gateway Protocol (IGP) system	Enables synchronization between BGP and your Interior Gateway Protocol (IGP) system. To enable the Cisco IOS software to advertise a network route without waiting for the IGP, deselect this option. Usually, a BGP speaker does not advertise a route to an external neighbor unless that route is local or exists in the IGP. By default, synchronization between BGP and the IGP is turned off to allow the Cisco IOS software to advertise a network route without waiting for route validation from the IGP. This feature allows routers and access servers within an autonomous system to have the route before BGP makes it available to other autonomous systems. Use synchronization if routers in the autonomous system do not speak BGP.
Redistribute iBGP into an IGP (use filtering to limit the number of prefixes that are redistributed)	(Optional) Configures iBGP redistribution into an interior gateway protocol (IGP), such as IS-IS or OSPF.
Administrative Route Distances	
External	Specifies the administrative distance for external BGP routes. Routes are external when learned from an external autonomous system. The range of values for this argument are from 1 to 255. The default value is 20.
Internal	Specifies administrative distance for internal BGP routes. Routes are internal when learned from peer in the local autonomous system. The range of values for this argument are from 1 to 255. The default value is 200.
Local	Specifies administrative distance for local BGP routes. Local routes are those networks listed with a network router configuration command, often as back doors, for the router or for the networks that is being redistributed from another process. The range of values for this argument are from 1 to 255. The default value is 200.
Forward packets over Multiple Paths	
Number of Paths	(Optional) Specify the maximum number of Border Gateway Protocol routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.
IBGP Number of Paths	(Optional) Specify the maximum number of parallel internal Border Gateway Protocol (iBGP) routes that can be installed in a routing table. The range of values are from 1 to 8. The default value is 1.

Add/Edit Aggregate Address Dialog Box

Use the Add/Edit Aggregate Address dialog box to define the aggregation of specific routes into one route.

Navigation Path

You can access the Add/Edit Aggregate Address dialog box from the IPv6 Family Tab, page 56-20.

Click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv6 Family - General Tab, page 56-21
- Add/Edit Neighbor Dialog Box, page 56-24
- Add/Edit Network Dialog Box, page 56-29
- Add/Edit Redistribution Dialog Box, page 56-30
- Add/Edit Route Injection Dialog Box, page 56-31

Field Reference

Table 56-12 Add/Edit Aggregate Address Dialog Box

Element	Description
Aggregate Timer	Specify a value for the aggregate timer (in seconds). Valid values are 0 or any value between 6 and 60. This specifies the interval at which the routes will be aggregated. The default value is 30 seconds.
Network	Enter an IPv6 address, or enter or Select the desired Network/Hosts objects.
Attribute Map	(Optional) Enter or Select the route map used to set the attribute of the aggregate route. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Advertise Map	(Optional) Enter or Select the route map used to select the routes to create AS_SET origin communities. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Suppress Map	(Optional) Enter or Select the route map used to select the routes to be suppressed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Generate AS Set Path Information	Enables generation of autonomous system set path information.

Table 56-12 Add/Edit Aggregate Address Dialog Box (continued)

Element	Description
Filter all more-specific routes from updates	Filters all more-specific routes from updates.

Add/Edit Neighbor Dialog Box

Use the Add/Edit Neighbor dialog box to define BGP neighbors and neighbor settings.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the IPv6 Family Tab, page 56-20.

Click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv6 Family - General Tab, page 56-21
- Add/Edit Aggregate Address Dialog Box, page 56-22
- Add/Edit Network Dialog Box, page 56-29
- Add/Edit Redistribution Dialog Box, page 56-30
- Add/Edit Route Injection Dialog Box, page 56-31

Field Reference

Table 56-13 Add/Edit Neighbor Dialog Box

Element	Description
General	
IP Address	Enter the BGP neighbor IPv6 address in one of the following formats: • IPv6 address • <IPv6 address>%<Interface name> This IPv6 address is added to the BGP neighbor table.
Remote AS	Enter the autonomous system to which the BGP neighbor belongs.
Enable Address Family	(Optional) Enables communication with the BGP neighbor.
Shutdown neighbor administratively	(Optional) Disable a neighbor or peer group.
Description	(Optional) Enter a description for the BGP neighbor.
fall-over BFD	(Optional) Enables BFD support for fall-over for the BGP neighbor.
BFD-Hop	(Optional) Specify if there is a single IP hop or multiple IP hops between a BFD source and destination.

Table 56-13 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Filtering	
Filter routes using route map	(Optional) Enter or Select the appropriate incoming or outgoing route maps to apply a route map to incoming or outgoing routes. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Filter routes using a Prefix list	(Optional) Enter or Select the appropriate incoming or outgoing prefix list to distribute BGP neighbor information. **Tip** Click **Select** to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box, page 56-146.
Filter routes using AS Path filter	(Optional) Enter or Select the appropriate incoming or outgoing AS path filter to distribute BGP neighbor information. **Tip** Click **Select** to open the AS Path Object Selector from which you can select an AS path object. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes, page 56-151.
Limit the number of prefixes allowed from the neighbor	(Optional) Select to control the number of prefixes that can be received from a neighbor. • Enter the maximum number of prefixes allowed from a specific neighbor in the **Maximum Prefixes** field. • Enter the percentage (of maximum) at which the router starts to generate a warning message in the **Threshold Level** field. Valid values are integers between 1 and 100. The default value is 75. • (Optional) Check the **Control prefixes received from the peer** check box to specify additional controls for the prefixes received from a peer. Do one of the following: – Select **Terminate peering when prefix limit is exceeded** to stop the BGP neighbor when the prefix limit is reached. Specify the interval after which the BGP neighbor will restart in the Restart interval field. – Select **Give only warning message when prefix limit is exceeded** to generate a log message when the maximum prefix limit is exceeded. Here, the BGP neighbor will not be terminated.
Routes	
Advertisement Interval	Enter the minimum interval (in seconds) between the sending of BGP routing updates. Valid values are between 0 and 600.

Table 56-13 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Remove private AS numbers from outbound routing updates	(Optional) Excludes the private AS numbers from being advertised on outbound routes.
Generate Default route	(Optional) Select to allow the local router to send the default route 0.0.0.0 to a neighbor to use as a default route. Enter or Select the route map that allows the route 0.0.0.0 to be injected conditionally in the Route map field. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Conditionally Advertised Routes	(Optional) To add or edit conditionally advertised routes, click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button. In the Add/Edit Advertised Route dialog box, do the following: • Click **Select** to open the Route Map Object Selector from which you can select a route map that will be advertised if the conditions of the exist map or the non-exist map are met. For more information about route maps, see Understanding Route Map Objects, page 56-132 • Do one of the following: – Select Set Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised. – Select Non-Exist Map and choose a route map from the Route Map Object Selector. This route map will be compared with the routes in the BGP table, to determine whether or not the advertise map route is advertised.
Timers	
Set timers for the BGP peer	(Optional) Select to set the keepalive frequency, hold time and minimum hold time.
Keepalive Interval	Enter the frequency (in seconds) with which the ASA sends keepalive messages to the neighbor. Valid values are between 0 and 65535. The default value is 60 seconds.
Hold Time	Enter the interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 180 seconds.
Min Hold Time	(Optional) Enter the minimum interval (in seconds) after not receiving a keepalive message that the ASA declares a peer dead. Valid values are between 0 and 65535. The default value is 0 seconds.

Table 56-13 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Advanced	
Enable Authentication	(Optional) Select to enable MD5 authentication on a TCP connection between two BGP peers. • Choose an encryption type from the Enable Encryption drop-down list. • Enter a password in the Password field. Reenter the password in the Confirm field. The password is case-sensitive and can be up to 25 characters long when the service password-encryption command is enabled and up to 81 characters long when the service password-encryption command is not enabled. The first character cannot be a number. The string can contain any alphanumeric characters, including spaces. **Note** You cannot specify a password in the format number-space-anything. The space after the number can cause authentication to fail.
Send Community attribute to this neighbor	(Optional) Specifies that communities attributes should be sent to the BGP neighbor.
Use ASA as next hop for neighbor	(Optional) Select to configure the router as the next-hop for a BGP speaking neighbor or peer group.
Disable connection verification	(Optional) Select to disable the connection verification process for eBGP peering sessions that are reachable by a single hop but are configured on a loopback interface or otherwise configured with a non-directly connected IPv6 address. This command is required only when the neighbor ebgp-multihop command is configured with a TTL value of 1. The address of the single-hop eBGP peer must be reachable. The neighbor update-source command must be configured to allow the BGP routing process to use the loopback interface for the peering session. When deselected (default), a BGP routing process will verify the connection of single-hop eBGP peering session (TTL=254) to determine if the eBGP peer is directly connected to the same network segment by default. If the peer is not directly connected to same network segment, connection verification will prevent the peering session from being established.
Allow connections with neighbor that is not directly connected	Select to accept and attempt BGP connections to external peers residing on networks that are not directly connected. (Optional) Enter the time-to-live in the TTL hops field. Valid values are between 1 and 255. **Note** This feature should be used only under the guidance of Cisco technical support staff. To prevent the creation of loops through oscillating routes, the multihop will not be established if the only route to the multihop peer is the default route (0.0.0.0).

Table 56-13 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Limit number of TTL hops to neighbor	Select this option to secure a BGP peering session. Enter the maximum number of hops that separate eBGP peers in the TTL hops field. Valid values are between 1 and 254. This feature provides a lightweight security mechanism to protect BGP peering sessions from CPU utilization-based attacks. These types of attacks are typically brute force Denial of Service (DoS) attacks that attempt to disable the network by flooding the network with IP packets that contain forged source and destination IP addresses in the packet headers. This feature leverages designed behavior of IP packets by accepting only IP packets with a TTL count that is equal to or greater than the locally configured value. Accurately forging the TTL count in an IP packet is generally considered to be impossible. Accurately forging a packet to match the TTL count from a trusted peer is not possible without internal access to the source or destination network. This feature should be configured on each participating router. It secures the BGP session in the incoming direction only and has no effect on outgoing IP packets or the remote router. When this feature is enabled, BGP will establish or maintain a session only if the TTL value in the IP packet header is equal to or greater than the TTL value configured for the peering session. This feature has no effect on the BGP peering session, and the peering session can still expire if keepalive packets are not received. If the TTL value in a received packet is less than the locally configured value, the packet is silently discarded and no Internet Control Message Protocol (ICMP) message is generated. This is designed behavior; a response to a forged packet is not necessary. To maximize the effectiveness of this feature, the hop-count value should be strictly configured to match the number of hops between the local and external network. However, you should also take path variation into account when configuring this feature for a multihop peering session. The following restrictions apply to the configuration of this command: • This feature is not supported for internal BGP (iBGP) peers. • The effectiveness of this feature is reduced in large-diameter multihop peerings. In the event of a CPU utilization-based attack against a BGP router that is configured for large-diameter peering, you may still need to shut down the affected peering sessions to handle the attack. • This feature is not effective against attacks from a peer that has been compromised inside of your network. This restriction also includes peers that are on the network segment between the source and destination network.
Use TCP Path MTU Discovery	(Optional) Select to enable a TCP transport session for a BGP session.

Table 56-13 Add/Edit Neighbor Dialog Box (continued)

Element	Description
TCP transport mode	Choose the TCP connection mode from the drop-down list. Options are Default, Active, or Passive.
Weight	(Optional) Enter a weight for the BGP neighbor connection.
BGP Version	Choose the BGP version that the ASA will accept from the drop-down list. The version can be set to 4-Only to force the software to use only Version 4 with the specified neighbor. The default is to use Version 4 and dynamically negotiate down to Version 2 if requested.
Migration **Note** This customization should only be used for AS migration, and should be removed after the transition has been completed. The procedure should be attempted only by an experienced network operator. Routing loops can be created through improper configuration.	
Customize the AS number for routes received from the neighbor	(Optional) Select to customize the AS_PATH attribute for routes received from an eBGP neighbor.
Local AS Number	Enter the local autonomous system number. Valid values are any valid autonomous system number from 1 to 4294967295 or 1.0 to 65535.65535.
Do not prepend local AS number to routes received from neighbor	(Optional) Select to prevent the local AS number from being prepended to any routes received from eBGP peer.
Replace real AS number with local AS number in routes received from neighbor	(Optional) Select to replace the real autonomous system number with the local autonomous system number in the eBGP updates. The autonomous system number from the local BGP routing process is not prepended.
Accept either real AS number or local AS number in routes received from neighbor	(Optional) Configures the eBGP neighbor to establish a peering session using the real autonomous system number (from the local BGP routing process) or by using the local autonomous system number.

Add/Edit Network Dialog Box

Use the Add/Edit Network dialog box to define the networks to be advertised by the BGP routing process.

Navigation Path

You can access the Add/Edit Network dialog box from the IPv6 Family Tab, page 56-20.

Click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv6 Family - General Tab, page 56-21
- Add/Edit Aggregate Address Dialog Box, page 56-22
- Add/Edit Neighbor Dialog Box, page 56-24

- Add/Edit Redistribution Dialog Box, page 56-30
- Add/Edit Route Injection Dialog Box, page 56-31

Field Reference

Table 56-14 Add/Edit Network Dialog Box

Element	Description
Prefix Name	(Optional) Enter a name for the prefix and enable the DHCpv6 Prefix Delegation client. Valid values are a string not exceeding 200 characters.
Network	Specifies the network to be advertised by the BGP routing processes.
Route Map	(Optional) Enter or Select a route map that should be examined to filter the networks to be advertised. If not specified, all networks are redistributed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.

Add/Edit Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to define the conditions for redistributing routes from another routing domain into BGP.

Navigation Path

You can access the Add/Edit Redistribution dialog box from the IPv6 Family Tab, page 56-20.

Click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv6 Family - General Tab, page 56-21
- Add/Edit Aggregate Address Dialog Box, page 56-22
- Add/Edit Neighbor Dialog Box, page 56-24
- Add/Edit Network Dialog Box, page 56-29
- Add/Edit Route Injection Dialog Box, page 56-31

Field Reference

Table 56-15 Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the protocol from which you want to redistribute routes into the BGP domain from the Source Protocol drop-down list.

Table 56-15 Add/Edit Redistribution Dialog Box (continued)

Element	Description
Process ID	Enter the identifier for the routing process. Applies to EIGRP and OSPF routing protocols.
Metric	(Optional) Enter a metric for the redistributed route.
Route Map	Enter or Select a route map that should be examined to filter the networks to be redistributed. If not specified, all networks are redistributed. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Include Connected	To include connected routes in the redistribution, check the Include Connected check box. This option is enabled only when OSPF is chosen as the Source Protocol.
Match	The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol. • Internal • External 1 • External 2 • NSSA External 1 • NSSA External 2

Add/Edit Route Injection Dialog Box

Use the Add/Edit Route Injection dialog box to define the routes to be conditionally injected into the BGP routing table.

Navigation Path

You can access the Add/Edit Route Injection dialog box from the IPv6 Family Tab, page 56-20.

Click the **Add Row (+)**button, or select a row in the table and click the **Edit Row(pencil)**button.

Related Topics

- Configuring BGP, page 56-2
- About BGP, page 56-3
- IPv6 Family - General Tab, page 56-21
- Add/Edit Aggregate Address Dialog Box, page 56-22
- Add/Edit Neighbor Dialog Box, page 56-24
- Add/Edit Network Dialog Box, page 56-29
- Add/Edit Redistribution Dialog Box, page 56-30

Field Reference

Table 56-16 Add/Edit Route Injection Dialog Box

Element	Description
Inject Map	Enter or Select the route map that specifies the prefixes to inject into the local BGP routing table. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Exist Map	Enter or Select the route map containing the prefixes that the BGP speaker will track. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Injected routes will inherit the attributes of the aggregate route	Configures the injected route to inherit attributes of the aggregate route.

Configuring EIGRP

The EIGRP page provides six tabbed panels for configuring Enhanced Interior Gateway Routing Protocol (EIGRP) routing on a firewall device. The following topics provide detailed information about enabling and configuring EIGRP:

- About EIGRP, page 56-33
- EIGRP Advanced Dialog Box, page 56-34
- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Neighbors Tab, page 56-41
- Redistribution Tab, page 56-42
- Summary Address Tab, page 56-45
- Interfaces Tab, page 56-47

Navigation Path

- (Device view) Select **Platform > Routing > EIGRP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > EIGRP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Field Reference

Table 56-17 EIGRP Page

Element	Description
Enable EIGRP	Check this box to enable the EIGRP routing process.
AS Number	Enter the autonomous system (AS) number for the EIGRP process. The AS number can be from 1 to 65535.
Advanced button	Opens the EIGRP Advanced Dialog Box, page 56-34, in which you can configure additional EIGRP process settings, such as the router ID, stub routing, and adjacency changes.
Setup tab	Use the Setup tab to configure the networks used by the EIGRP routing process, passive interfaces, default route information, administrative distances, and default metrics. For more information, see Setup Tab, page 56-36.
Filter Rules tab	Use the Filter Rules tab to define filter rules that let you control which routes are accepted or advertised by the EIGRP routing process. For more information, see Filter Rules Tab, page 56-39.
Neighbors tab	Use the Neighbors tab to manually define EIGRP neighbors. For more information, see Neighbors Tab, page 56-41.
Redistribution tab	Use the Redistribution tab to define the rules for redistributing routes from other routing protocols to the EIGRP routing process. For more information, see Redistribution Tab, page 56-42.
Summary Address tab	Use the Summary Address tab to create statically defined EIGRP summary addresses. For more information, see Summary Address Tab, page 56-45.
Interfaces tab	Use the Interfaces tab to configure interfaces for EIGRP. For more information, see Interfaces Tab, page 56-47.

About EIGRP

EIGRP is an enhanced version of IGRP developed by Cisco. Unlike IGRP and RIP, EIGRP does not send out periodic route updates. EIGRP updates are sent out only when the network topology changes. Key capabilities that distinguish EIGRP from other routing protocols include fast convergence, support for variable-length subnet mask, support for partial updates, and support for multiple network layer protocols.

A router running EIGRP stores all the neighbor routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found. Its support for variable-length subnet masks permits routes to be automatically summarized on a network number boundary. In addition, EIGRP can be configured to summarize on any bit boundary at any interface. EIGRP does not make periodic updates. Instead, it sends partial updates only when the metric for a route changes. Propagation of partial updates is automatically bounded so that only those routers that need the information are updated. As a result of these two capabilities, EIGRP consumes significantly less bandwidth than IGRP.

Neighbor discovery is the process that the ASA uses to dynamically learn of other routers on directly attached networks. EIGRP routers send out multicast hello packets to announce their presence on the network. When the ASA receives a hello packet from a new neighbor, it sends its topology table to the neighbor with an initialization bit set. When the neighbor receives the topology update with the initialization bit set, the neighbor sends its topology table back to the ASA.

The hello packets are sent out as multicast messages. No response is expected to a hello message. The exception to this is for statically defined neighbors. If you manually configure a neighbor, the hello messages sent to that neighbor are sent as unicast messages. Routing updates and acknowledgments are sent out as unicast messages.

Once this neighbor relationship is established, routing updates are not exchanged unless there is a change in the network topology. The neighbor relationship is maintained through the hello packets. Each hello packet received from a neighbor includes a hold time. This is the time in which the ASA can expect to receive a hello packet from that neighbor. If the ASA does not receive a hello packet from that neighbor within the hold time advertised by that neighbor, the ASA considers that neighbor to be unavailable.

The EIGRP protocol uses four key algorithm technologies, four key technologies, including neighbor discovery/recovery, Reliable Transport Protocol (RTP), and DUAL, which is important for route computations. DUAL saves all routes to a destination in the topology table, not just the least-cost route. The least-cost route is inserted into the routing table. The other routes remain in the topology table. If the main route fails, another route is chosen from the feasible successors. A successor is a neighboring router used for packet forwarding that has a least-cost path to a destination. The feasibility calculation guarantees that the path is not part of a routing loop.

If a feasible successor is not found in the topology table, a route recomputation must occur. During route recomputation, DUAL queries the EIGRP neighbors for a route, who in turn query their neighbors. Routers that do no have a feasible successor for the route return an unreachable message.

During route recomputation, DUAL marks the route as active. By default, the ASA waits for three minutes to receive a response from its neighbors. If the ASA does not receive a response from a neighbor, the route is marked as stuck-in-active. All routes in the topology table that point to the unresponsive neighbor as a feasibility successor are removed.

Note EIGRP neighbor relationships are not supported through the IPsec tunnel without a GRE tunnel.

Related Topics

- Configuring EIGRP, page 56-32

EIGRP Advanced Dialog Box

Use the EIGRP Advanced dialog box to configure settings such as the router ID, stub routing, and adjacency changes.

Navigation Path

You can access the EIGRP Advanced dialog box from the EIGRP page (see Configuring EIGRP, page 56-32).

Related Topics

- Configuring EIGRP, page 56-32

Field Reference

Table 56-18 EIGRP Advanced Dialog Box

Element	Description
Router ID	The router ID is used to identify the originating router for external routes. If an external route is received with the local router ID, the route is discarded. To prevent this, specify a global address for the router ID. A unique value should be configured for each EIGRP router. On a single device, choose **Automatic** or **IP Address**. (An address field appears when you choose IP Address.) If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field. On a device cluster, choose **Automatic** or **Cluster Pool**. (An IPv4 Pool object ID field appears when you choose Cluster Pool.) If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see Add or Edit IPv4 Pool Dialog Box, page 6-92.
Stub	You can enable, and configure the ASA as an EIGRP stub router. Stub routing decreases memory and processing requirements on the ASA. As a stub router, the ASA does not need to maintain a complete EIGRP routing table because it forwards all nonlocal traffic to a distribution router. Generally, the distribution router need not send anything more than a default route to the stub router. Only specified routes are propagated from the stub router to the distribution router. As a stub router, the ASA responds to all queries for summaries, connected routes, redistributed static routes, external routes, and internal routes with the message “inaccessible.” When the ASA is configured as a stub, it sends a special peer information packet to all neighboring routers to report its status as a stub router. Any neighbor that receives a packet informing it of the stub status will not query the stub router for any routes, and a router that has a stub peer will not query that peer. The stub router depends on the distribution router to send the correct updates to all peers. To enable the ASA as an EIGRP stub routing process, choose one or more of the following EIGRP stub routing processes: • Receive only—Configures the EIGRP stub routing process to receive route information from the neighbor routers but does not send route information to the neighbors. If this option is selected, you cannot select any of the other stub routing options. • Connected—Advertises connected routes. • Redistributed—Advertises redistributed routes. • Static—Advertises static routes. • Summary—Advertises summary routes.

Table 56-18 EIGRP Advanced Dialog Box (continued)

Element	Description
Adjacency Changes	These options specify the syslog messages sent when adjacency changes occur. • Log Neighbor Changes – enables the logging of EIGRP neighbor adjacency changes. This option is selected by default. • Log Neighbor Warnings – enables the logging of EIGRP neighbor warning messages. This option is selected by default. (Optional) The time interval (in seconds) between repeated neighbor warning messages. Valid values are from 1 to 65535. Repeated warnings are not logged if they occur during this interval.

Setup Tab

Use the Setup tab on the EIGRP page to configure the networks used by the EIGRP routing process, passive interfaces, default route information, administrative distances, and default metrics.

Navigation Path

You can access the Setup tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Neighbors Tab, page 56-41
- Redistribution Tab, page 56-42
- Summary Address Tab, page 56-45
- Interfaces Tab, page 56-47

Field Reference

Table 56-19 EIGRP - Setup Tab

Element	Description
Auto Summary	Check this box to enable automatic route summarization. Auto summary is enabled by default for ASA versions earlier than 9.2.1 and is disabled by default for ASA 9.2(1) and later. When enabled, the EIGRP routing process summarizes on network number boundaries. This can cause routing problems if you have noncontiguous networks. For example, if you have a router with the networks 192.168.1.0, 192.168.2.0, and 192.168.3.0 connected to it, and those networks all participate in EIGRP, the EIGRP routing process creates the summary address 192.168.0.0 for those routes. If an additional router is added to the network with the networks 192.168.10.0 and 192.168.11.0, and those networks participate in EIGRP, they will also be summarized as 192.168.0.0. To prevent the possibility of traffic being routed to the wrong location, you should disable automatic route summarization on the routers creating the conflicting summary addresses.
Networks	Enter the IP addresses of the networks to participate in the EIGRP routing process. **Tip** You can click **Select** to select the networks from a list of network/host objects.
Passive Interface	You can configure one or more interfaces as passive interfaces. In EIGRP, a passive interface does not send or receive routing updates. By deafult, all interfaces are enabled for active routing (sending and receiving routing updates) when routing is enabled for that interface. To configure passive interfaces, do one of the following: • To enable all interfaces for active routing (sending and receiving routing updates) when routing is enabled for that interface, select **None**. • To configure all interfaces as passive, select **All Interfaces**. • To configure specific interfaces as passive, select **Specified Interfaces** and then enter or select the interfaces that you want to make passive.

Table 56-19 EIGRP - Setup Tab (continued)

Element	Description
Default Route Information	You can control the sending and receiving of default route information in EIGRP updates. By default, default routes are sent and accepted. Configuring the ASA to disallow default information to be received causes the candidate default route bit to be blocked on received routes. Configuring the ASA to disallow default information to be sent disables the setting of the default route bit in advertised routes. • Accept Default Route Info—configures EIGRP to accept exterior default routing information. Optionally, you can specify a standard access list that define which networks are allowed and which are not when receiving default route information. • Send Default Route Info—configures EIGRP to advertise external routing information. Optionally, you can specify a standard access list that defines which networks are allowed and which are not when sending default route information.
Administrative Distance	Because every routing protocol has metrics based on algorithms that are different from the other routing protocols, it is not always possible to determine the “best path” for two routes to the same destination that were generated by different routing protocols. Administrative distance is a route parameter that the ASA uses to select the best path when there are two or more different routes to the same destination from two different routing protocols. If you have more than one routing protocol running on the ASA, you can use the distance eigrp command to adjust the default administrative distances of routes discovered by the EIGRP routing protocol in relation to the other routing protocols.: Internal Distance—Administrative distance for EIGRP internal routes. Internal routes are those that are learned from another entity within the same autonomous system. Valid values are from 1 to 255. The default value is 90. External Distance—Administrative distance for EIGRP external routes. External routes are those for which the best path is learned from a neighbor external to the autonomous system. Valid values are from 1 to 255.The default value is 170.

Table 56-19 EIGRP - Setup Tab (continued)

Element	Description
Default Metrics	You can define the default metrics for routes redistributed into the EIGRP routing process: • Bandwidth—the minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295. • Delay Time—the route delay in tens of microseconds. Valid values range from 0 to 4294967295. • Reliability—the likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability. • Loading—the effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded. • MTU—the smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.

Filter Rules Tab

The Filter Rules tab contains the Filter Rules table which displays the route filtering rules configured for the EIGRP routing process. Filter rules let you control which routes are accepted or advertised by the EIGRP routing process.

Navigation Path

You can access the Filter Rules tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Add/Edit EIGRP Filter Rule Dialog Box, page 56-40
- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Setup Tab, page 56-36
- Neighbors Tab, page 56-41
- Redistribution Tab, page 56-42
- Summary Address Tab, page 56-45
- Interfaces Tab, page 56-47

Field Reference

Table 56-20 EIGRP - Filter Rules Tab

Element	Description
Direction	The direction for the filter rule: • Inbound—The rule filters default route information from incoming EIGRP routing updates. • Outbound—The rule filters default route information from outgoing EIGRP routing updates.
Interface	(Optional) The interface to which the filter rule applies.
Protocol	The routing protocol being filtered: BGP, Connected, OSPF, RIP, or Static.
ACL	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.

Add/Edit EIGRP Filter Rule Dialog Box

Use the Add/Edit EIGRP Filter Rule dialog box to add new filter rules to the Filter Rules table or to modify an existing filter rule.

Navigation Path

You can access the Add/Edit EIGRP Filter Rule dialog box from the Filter Rules Tab, page 56-39.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Filter Rules Tab, page 56-39

Field Reference

Table 56-21 Add/Edit EIGRP Filter Rule Dialog Box

Element	Description
EIGRP Filter Direction	Specify the direction for the filter rule: • Inbound—The rule filters default route information from incoming EIGRP routing updates. • Outbound—The rule filters default route information from outgoing EIGRP routing updates.

Table 56-21 Add/Edit EIGRP Filter Rule Dialog Box (continued)

Element	Description
Type	Specify the type of filter rule: • (Optional) Interface—Specify the interface on which to apply the routing updates. Specifying an interface causes the access list to be applied only to routing updates for that interface. If no interface is specified, the access list will be applied to all updates. • (Optional) Routing Protocol—For outbound EIGRP routing updates, select the routing protocol for which you want to filter: BGP, Connected, OSPF, RIP, or Static. Routing Protocol ID—Enter the identifier for the routing process. Applies to BGP and OSPF routing protocols.
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.

Neighbors Tab

The Neighbors tab contains the Neighbors table, through which you can define static neighbors. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

Navigation Path

You can access the Neighbors tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Add/Edit EIGRP Neighbor Dialog Box, page 56-42
- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Redistribution Tab, page 56-42
- Summary Address Tab, page 56-45
- Interfaces Tab, page 56-47

Field Reference

Table 56-22 EIGRP - Neighbors Tab

Element	Description
Interface	The interface through which the neighbor is available.
Neighbor	The IP address of the static neighbor.

Add/Edit EIGRP Neighbor Dialog Box

EIGRP hello packets are sent as multicast packets. If an EIGRP neighbor is located across a non broadcast network, such as a tunnel, you must manually define that neighbor. When you manually define an EIGRP neighbor, hello packets are sent to that neighbor as unicast messages.

Note Configuring the passive-interface command for an interface suppresses all incoming and outgoing routing updates and hello messages on that interface. EIGRP neighbor adjacencies cannot be established or maintained over an interface that is configured as passive.

Use the Add/Edit EIGRP Neighbor dialog box to define a static neighbor or change information for an existing static neighbor.

Navigation Path

You can access the Add/Edit EIGRP Neighbor dialog box from the Neighbors Tab, page 56-41.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Neighbors Tab, page 56-41

Field Reference

Table 56-23 Add/Edit EIGRP Neighbor Dialog Box

Element	Description
Interface	The interface through which the neighbor is available. **Tip** You can click **Select** to select the interface from a list of interface objects.
Neighbor	The IP address of the static neighbor. **Tip** You can click **Select** to select the neighbor from a list of host objects.

Redistribution Tab

Use the Redistribution tab to define the rules for redistributing routes from other routing protocols to the EIGRP routing process.

Navigation Path

You can access the Redistribution tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Add/Edit EIGRP Redistribution Dialog Box, page 56-44
- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33

- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Neighbors Tab, page 56-41
- Summary Address Tab, page 56-45
- Interfaces Tab, page 56-47

Field Reference

Table 56-24 EIGRP - Redistribution Tab

Element	Description
Protocol	The source protocol from which the routes are being redistributed: • BGP—Redistribute routes discovered by the BGP routing process to EIGRP. • RIP—Redistributes routes discovered by the RIP routing process to EIGRP. • Static—Redistributes static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • OSPF—Redistributes routes discovered by the OSPF routing process to EIGRP. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or BGP routes.
ID	The autonomous system (AS) number for the BGP or OSPF routing process.
Bandwidth	The minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295.
Delay Time	The route delay in tens of microseconds. Valid values range from 0 to 4294967295.
Reliability	The likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability.
Loading	The effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded.
MTU	The smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.
Route Map	The name of the route map object to apply to the redistribution entry.

Add/Edit EIGRP Redistribution Dialog Box

Use the Add/Edit Redistribution dialog box to add a redistribution rule or to edit an existing redistribution rule in the Redistribution table.

Navigation Path

You can access the Add/Edit EIGRP Redistribution dialog box from the Redistribution Tab, page 56-42.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Redistribution Tab, page 56-42

Field Reference

Table 56-25 Add/Edit EIGRP Redistribution Dialog Box

Element	Description
Protocol	Select the source protocol from which the routes are being redistributed. You can choose one of the following options: • BGP—Redistribute routes discovered by the BGP routing process to EIGRP. • RIP—Redistributes routes discovered by the RIP routing process to EIGRP. • Static—Redistributes static routes to the EIGRP routing process. Static routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the EIGRP routing process. Connected routes that fall within the scope of a network statement are automatically redistributed into EIGRP; you do not need to define a redistribution rule for them. • OSPF—Redistributes routes discovered by the OSPF routing process to EIGRP. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, or BGP routes.
Routing Process ID	The autonomous system (AS) number for the BGP or OSPF routing process.

Table 56-25 Add/Edit EIGRP Redistribution Dialog Box (continued)

Element	Description
Optional Metrics	You can define the following metrics for routes redistributed into the EIGRP routing process: • Bandwidth—the minimum bandwidth of the route in kilobits per second. Valid values range from 1 to 4294967295. • Delay Time—the route delay in tens of microseconds. Valid values range from 0 to 4294967295. • Reliability—the likelihood of successful packet transmission expressed as a number 0 through 255. The value 255 indicates 100 percent reliability; 0 means no reliability. • Loading—the effective bandwidth of the route. Valid values range from 1 to 255; 255 indicates 100 percent loaded. • MTU—the smallest allowed value for the maximum transmission unit of the path. Valid values range from 1 to 65535.
Route Map	Enter or Select a route map object to define which routes are redistributed into the EIGRP routing process. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Optional OSPF Redistribution	If you have chosen OSPF as the Route Type, choose the conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions: • Internal—The route is internal to a specific AS. • External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes. • External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes. • NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes. • NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Summary Address Tab

Use the Summary Address tab to configure a summary for EIGRP on a specific interface. You can configure summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Navigation Path

You can access the Summary Address tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Add/Edit EIGRP Summary Address Dialog Box, page 56-46
- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Neighbors Tab, page 56-41
- Redistribution Tab, page 56-42
- Interfaces Tab, page 56-47

Field Reference

Table 56-26 EIGRP - Summary Address Tab

Element	Description
Interface	The interface from which the summary address is advertised.
Network	The IP address and network mask of the summary address.
Administrative Distance	The administrative distance of the summary route.

Add/Edit EIGRP Summary Address Dialog Box

Use the Add/Edit EIGRP Summary Address dialog box to add new entries or to modify existing entries in the Summary Address table. You can configure summary addresses on a per-interface basis. You need to manually define summary addresses if you want to create summary addresses that do not occur at a network number boundary or if you want to use summary addresses on an ASA with automatic route summarization disabled. If any more specific routes are in the routing table, EIGRP will advertise the summary address out the interface with a metric equal to the minimum of all more specific routes.

Navigation Path

You can access the Add/Edit EIGRP Summary Address dialog box from the Summary Address Tab, page 56-45.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Summary Address Tab, page 56-45

Field Reference

Table 56-27 Add/Edit EIGRP Summary Address Dialog Box

Element	Description
Interface	The interface from which the summary address is advertised. **Tip** You can click **Select** to select the interface from a list of interface objects.
Networks	The IP address and network mask of the summary address. **Tip** You can click **Select** to select the network from a list of network objects.
Administrative Distance	(Optional) The administrative distance of the summary route. Valid values are from 1 to 255. The default value is 5.

Interfaces Tab

Use the Interfaces tab to configure interface-specific EIGRP routing properties.

Navigation Path

You can access the Interfaces tab from the EIGRP Page; see Configuring EIGRP, page 56-32 for more information.

Related Topics

- Add/Edit EIGRP Interface Dialog Box, page 56-48
- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Setup Tab, page 56-36
- Filter Rules Tab, page 56-39
- Neighbors Tab, page 56-41
- Redistribution Tab, page 56-42
- Summary Address Tab, page 56-45

Field Reference

Table 56-28 EIGRP - Interfaces Tab

Element	Description
Interface	The name of the interface to which the configuration applies.
Hello Interval	The interval, in seconds, between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Hold Time	The hold time advertised by the ASA in EIGRP hello packets. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
Split Horizon	Whether EIGRP split-horizon is enabled (true) or disabled (false) on an interface.

Table 56-28 EIGRP - Interfaces Tab (continued)

Element	Description
Delay	The delay time in tens of microseconds. Valid values are from 1 to 16777215. This option is not supported for devices in multi-context mode.
Key ID	The ID of the key used to authenticate EIGRP updates.

Add/Edit EIGRP Interface Dialog Box

Use the Add/Edit EIGRP Interface dialog box to configure interface-specific EIGRP routing parameters.

Navigation Path

You can access the Add/Edit EIGRP Interface dialog box from the Interfaces Tab, page 56-47.

Related Topics

- Configuring EIGRP, page 56-32
- About EIGRP, page 56-33
- Interfaces Tab, page 56-47

Field Reference

Table 56-29 Add/Edit EIGRP Interface Dialog Box

Element	Description
Interface	The name of the interface to which the configuration applies.
Hello Interval	The interval, in seconds, between EIGRP hello packets sent on an interface. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Hold Time	The hold time advertised by the ASA in EIGRP hello packets. Valid values range from 1 to 65535 seconds. The default value is 15 seconds.
Split Horizon	Enable/disable EIGRP split-horizon on an interface.
Delay Time	The delay time in tens of microseconds. Valid values are from 1 to 16777215. This option is not supported for devices in multi-context mode and will be disabled.
Enable MD5 Authentication	Enables MD5 authentication of EIGRP packets.
Key Type	Select Clear Text to indicate that the key you will be entering is in clear text. Select Encrypted to indicate that the key you will be entering is already encrypted.
Key ID and Key	Specify the key to authenticate EIGRP updates: • Key ID—Enter a numerical key identifier. Valid values range from 0to 255. • Key—An alphanumeric character string of up to 16 bytes. • Confirm—Re-enter the key.

Configuring ISIS

The ISIS page provides nine tabbed panels for configuring ISIS (Intermediate System-to-Intermediate System) routing on a firewall device. ISIS routing protocol is supported from Security Manager version 4.11 for ASA devices running the software version 9.6(1) or later. The following topics provide detailed information about enabling and configuring ISIS:

Navigation Path

- (Device view) Select **Platform > Routing > ISIS** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > ISIS** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Select Enable ISIS to enable Intermediate System-to-Intermediate System routing protocol on the selected ASA device.

About ISIS

Intermediate System-to-Intermediate System (ISIS) routing protocol is a link-state Interior Gateway Protocol (IGP). Link-state protocols are characterized by the propagation of the information required to build a complete network connectivity map on each participating router. That map is then used to calculate the shortest path to destinations. The IOS ISIS implementation supports CLNP, IPv4, and IPv6.

A routing domain may be divided into one or more sub-domains. Each sub-domain is referred to as an area and is assigned an area address. Routing within an area is referred to as Level-1 routing. Routing between Level-1 areas is referred to as Level-2 routing. A router in OSI terminology is referred to as an Intermediate System (IS). An IS may operate at Level 1, Level 2, or both. ISs that operate at Level 1 exchange routing information with other Level-1 ISs in the same area. ISs that operate at Level 2 exchange routing information with other Level-2 routers regardless of whether they are in the same Level-1 area. The set of Level-2 routers and the links that interconnect them form the Level-2 sub-domain, which must not be partitioned in order for routing to work properly.

General Tab

Field Reference

Table 56-30 ISIS General Tab

Element	Description
Shutdown Protocol	To enable or disable ISIS protocol on an interface so that it cannot form any adjacency on any interface and will clear the ISIS link-state packet (LSP) database.
Use Dynamic Hostname	To enable or disable ISIS dynamic hostname capability on the router.
Do not pad LAN Hello PDUs	Check this box to prevent ISIS from padding LAN hello PDUs. ISIS hellos are padded to the full maximum transmission unit (MTU) size. This allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces. You can disable hello padding to avoid wasting network bandwidth in case the MTU of both interfaces is the same or in the case of translational bridging.

Table 56-30 ISIS General Tab (continued)

Element	Description
Advertise Passive Only	To configure ISIS to advertise only prefixes that belong to passive interfaces. It excludes IP prefixes of connected networks from link-state packet (LSP) advertisements, which reduces ISIS convergence time.
Act as	Choose whether to have your ASA act as station router, an area router, or both by clicking the appropriate radio button.
Topology Priority	To configure the priority of designated routers. Enter a value between 0 to 127. The default value is 64.
Route Priority Tag	Enter a tag that indicates the ASA device's route priority.
Conditionally Advertise as L2	Click ... to select from the available Route Map objects.
Adjacency	
Log Changes in Adjacency	To configure the ASA to send a syslog message when an ISIS neighbor goes up or down.
Include changes generated by non IIH event	To enable logging of adjacency changes including changes generated by non-IIH events.

IPv4 Family Tab

The IPv4 Family Tab has three tabbed panels for configuring IPv4 ISIS.

IPv4 Family Tab—General Tab

Field Reference

Table 56-31 ISIS IPv4 Family Tab—General Tab

Element	Description
Perform Adjacency Check	Check the 'Perform adjacency check' check box for the router to check on nearby IS routers.
Distance	
Administrative Distance	In the Administrative Distance field, enter a distance assigned to routes discovered by IS-IS protocol. Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 115.
Maximum No. of Forward Paths	Enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8, the default is 4.
Distribute Default Route	Check the Distribute default route check box to configure an IS routing process to distribute a default route, and then choose the default route from the Route Map Object selector.
ISIS Metrics	
Global ISIS Metric Level 1	Enter a number specifying the metric. The range depends on the TLV Style that you select. The default is 10. • If you select Use old style of TLVs with narrow metric, the range is 1 to 63. • If you select Use new style of TLVs to carry wider metric, the range is 1 to 16777214. • If you select Send and accept both styles of TLVs during transition, the range is 1 to 16777214.
Global ISIS Metric Level 2	Enter a number specifying the metric. The range depends on the TLV Style that you select. The default is 10. • If you select Use old style of TLVs with narrow metric, the range is 1 to 63. • If you select Use new style of TLVs to carry wider metric, the range is 1 to 16777214. • If you select Send and accept both styles of TLVs during transition, the range is 1 to 16777214.

Table 56-31 ISIS IPv4 Family Tab—General Tab (continued)

Element	Description
TLV Style	Select one of the following Type, Length, and Values: • Use old style of TLVs with narrow metric • Use new style of TLVs to carry wider metric • Send and accept both styles of TLVs during transition
Accept both styles of TLVs during transition	If you selected one of the first two options in TLV Style, you can select this option,
Apply metric style to	Select one of the following: • Level 1 • Level 2 • Both The default is Level 1.

IPv4 Family Tab—SPF Tab

Field Reference

Table 56-32 ISIS IPv4 FamilyTab—SPF Tab

Element	Description
Shortest Path First	
Honour external metrics during SPF calculations	Check this check box to have the SPF calculations include external metrics.
Signal other routers to not use this router as an intermediate hop in their SPF calculations	Check this check box if you want to exclude this device, and configure the following:
Specify on-startup behavior	If you select this element you must choose one of the following options: • Advertise overself as overloaded until BGP has converged • Specify time to advertise overself as overloaded after reboot—Specify the time in the range of 5 to 86400 seconds.
Don’t advertise IP prefixes learned from other protocols when overload bit is set	Check this check box to exclude IP prefixes.
Don’t advertise IP prefixes learned from another ISIS level when overload bit is set	Check this check box to exclude IP prefixes.
Minimum interval between partial route calculations	
PRC Interval	Enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.

Table 56-32 ISIS IPv4 FamilyTab—SPF Tab (continued)

Element	Description
Initial wait for PRC	Enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
Minimum wait between first and second PRC	Enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.
Minimum interval between SPF calculations	
Configure parameters for level 1	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Configure parameters for level 2	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

IPv4 Family Tab—Redistribution Tab

Use the Add/Edit button to add a new Redistribution route or edit an existing row.

Field Reference

Table 56-33 ISIS IPv4 Family Tab—Redistribution Tab

Element	Description
Source Protocol	From the Source Protocol drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
Process ID	Enter the Process ID for the source protocol.
Route Level	From the Route Level drop-down list, choose Level-1, Level- 2, or Level 1-2.
Metric	In the Metric field, enter a metric for the redistributed route. The range is 1 to 4294967295.

Table 56-33 ISIS IPv4 Family Tab—Redistribution Tab (continued)

Element	Description
Metric Type	For the Metric Type, click the internal or external radio button.
ISIS Inter Area Route Levels	
Source ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Destination ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Distribution List	Select from the available Access Control List or add new.
Route Map	Choose a route map from the Route Map Object selector that should be examined to filter the networks to be redistributed, or click Add to add a new route map or edit an existing route map.
Match	Check one or more of the Match check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

IPv6 Family Tab

The IPv6 Family Tab has three tabbed panels for configuring ISIS for IPv6 addresses.

Check the Enable IPv6 Family checkbox if you want to enable IPv6 for ISIS.

IPv6 Family Tab—General Tab

Field Reference

Table 56-34 ISIS IPv6 Family Tab—General Tab

Element	Description
Perform Adjacency Check	Check the 'Perform adjacency check' check box for the router to check on nearby IS routers.
Distance	
Administrative Distance	In the Administrative Distance field, enter a distance assigned to routes discovered by ISIS protocol. Administrative distance is a parameter used to compare routes among different routing protocols. In general, the higher the value, the lower the trust rating. And administrative distance of 255 means that the routing information source cannot be trusted at all and should be ignored. The range is 1 to 255. The default is 115.
Maximum No. of Forward Paths	Enter the maximum number of IS routes that can be installed in a routing table. The range is 1 to 8. The default is 4.
Distribute Default Route	Check the Distribute default route check box to configure an IS routing process to distribute a default route, and then choose the default route from the Route Map Object selector.

IPv6 Family Tab—SPF Tab

Field Reference

Table 56-35 ISIS IPv6 FamilyTab—SPF Tab

Element	Description
Shortest Path First	
Signal other routers to not use this router as an intermediate hop in their SPF calculations	Check this check box if you want to exclude this device, and configure the following:
Specify on-startup behavior	If you select this element you must choose one of the following options: • Advertise overself as overloaded until BGP has converged • Specify time to advertise overself as overloaded after reboot—Specify the time in the range of 5 to 86400 seconds.

Table 56-35 ISIS IPv6 FamilyTab—SPF Tab (continued)

Element	Description
Don't advertise IP prefixes learned from other protocols when overload bit is set	Check this check box to exclude IP prefixes.
Don't advertise IP prefixes learned from another ISIS level when overload bit is set	Check this check box to exclude IP prefixes.
Minimum interval between partial route calculations	
PRC Interval	Enter an amount of time for the router to wait between partial route calculations (PRCs). The range is 1 to 120 seconds. The default is 5 seconds.
Initial wait for PRC	Enter the initial PRC calculation delay (in milliseconds) after a topology change. The range is 1 to 120.000 milliseconds. The default is 2000 milliseconds.
Minimum wait between first and second PRC	Enter the amount of time in milliseconds that you want the router to wait between PRCs. The range is 1 to 120,000 milliseconds. The default is 5000 milliseconds.
Minimum interval between SPF calculations	
Configure parameters for level 1	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.
Configure parameters for level 2	
SPF calculation interval	Enter an amount of time for the router to wait between SPF calculations. The range is 1 to 120 seconds. The default is 10 seconds.
Initial wait for SPF calculation	Enter the amount of time for the router to wait for an SPF calculation. The range is 1 to 120.000 milliseconds. The default is 5500 milliseconds.
Minimum wait between first and second SPF calculation	Enter the amount of time in milliseconds that you want the router to wait between SPF calculations. The range is 1 to 120,000 milliseconds. The default is 5500 milliseconds.

IPv6 Family Tab—Redistribution Tab

Use the Add/Edit button to add or edit Redistribution routes.

Field Reference

Table 56-36 ISIS IPv6 Family Tab—Redistribution Tab

Element	Description
Source Protocol	From the Source Protocol drop-down list, choose the protocol (BGP, Connected, EIGRP, OSPF, RIP, or Static) from which you want to redistribute routes into the ISIS domain.
Process ID	Enter the Process ID for the source protocol.
Route Level	From the Route Level drop-down list, choose Level-1, Level- 2, or Level 1-2.
Metric	In the Metric field, enter a metric for the redistributed route. The range is 1 to 4294967295.
Metric Type	For the Metric Type, click the internal or external radio button.
ISIS Inter Area Route Levels	
Source ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Destination ISIS Level	Select Level 1 or Level 2. The default is Level 1.
Distribution List	Select from the available Access Control List or add new.
Route Map	Choose a route map from the Route Map Object selector that should be examined to filter the networks to be redistributed, or click Add to add a new route map or edit an existing route map.
Match	Check one or more of the Match check boxes -Internal, External 1, External 2, NSSA External 1, and NSSA External 2 check boxes to redistribute routes from an OSPF network.

IPv6 Family Tab—Summary Prefix

You must configure at least one Network Entity Title entry to proceed.

See Network Entity Title Tab, page 56-61 for more information.

Use the Add/Edit button to add or edit Summary Prefix.

Field Reference

Table 56-37 ISIS IPv6 Family Tab—Summary Prefix Tab

Element	Description
IPv6 Summary Prefix	IPv6 prefix in the form X.X.X.X.::X/0-128

Table 56-37 ISIS IPv6 Family Tab—Summary Prefix Tab (continued)

Element	Description
Apply Summary Prefix into	Select Level 1, Level 2, or Both. Level 1: Only routes redistributed into Level 1 are summarized with the configured address and mask value. Level 2: Routes learned by Level 1 routing are summarized into the Level 2 backbone with the configured address and mask value. Redistributed routes into Level 2 ISIS are also summarized. Both: Summary routes are applied when redistributing routes into Level 1 and Level2 ISIS and when Level 2 ISIS advertises Level 1 routes as reachable in it area.

Authentication Tab

Field Reference

Table 56-38 ISIS Authentication Tab

Element	Description
Configure authentication parameter for level 1	
Type	Select a Type from the drop-down list.
Key	Enter the key to authenticate ISIS updates. The key can include up to 16 characters.
Confirm	Confirm the key.
Send only	Click Enable or Disable depending on whether you want Send Only enabled.
Mode	Choose the authentication mode by clicking either the Disabled, MD5, or Clear Text radio buttons.
Area password	Enter the Area password and confirm the same in the next textbox.
Configure authentication parameter for level 2	
Type	Select a Type from the drop-down list.
Key	Enter the key to authenticate ISIS updates. The key can include up to 16 characters.
Confirm	Confirm the key.
Send only	Click Enable or Disable depending on whether you want Send Only enabled.
Mode	Choose the authentication mode by clicking either the Disabled, MD5, or Clear Text radio buttons.
Domain password	Enter the Domain password and confirm the same.

Link State Packet Tab

Field Reference

Table 56-39 ISIS Link State Packet Tab

Element	Description
Ignore LSP Errors	Check the Ignore LSP Errors check box to allow the ASA to ignore LSP packets that are received with internal checksum errors rather than purging the LSPs.
Flood LSPs before running SPF	Check this box to fast-flood and fill LSPs before running SPF. If you select this option, enter the number of LSPs to be flooded in the range of 1 to 15. This parameter sends a specified number of LSPs from the ASA. If no LSP number is specified, the default of 5 is used. The LSPs invoke SPF before running SPF. Cisco recommends that you enable fast flooding, because then you speed up the LSP flooding process, which improves overall network convergence time.The default value is 5.
Suppress IP prefixes	To suppress IP prefixes, check the Suppress IP prefixes check box, and then check one of the following. In networks where there is no limit placed on the number of redistributed routes into IS-IS, it is possible that the LSP can become full and routes will be dropped. Use these options to control which routes are suppressed when the PDU becomes full.
Don't advertise IP prefixes learned from another ISIS level when ran out of LSP fragments	Suppresses any routes coming from another level. For example, if the Level-2 LSP becomes full, routes from Level 1 are suppressed.
Don't advertise IP prefixes learned from other protocols when ran out of LSP fragments	Suppresses any redistributed routes on the ASA.
LSP General Interval	
LSP Interval Parameters for level 1	
LSP Calculation Interval	Enter the interval of time in seconds between transmission of each LSP. The range is 1-120 seconds. The default is 5. The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.
Initial wait for LSP calculation	Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.
Minimum wait between first and second	Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.

Table 56-39 ISIS Link State Packet Tab (continued)

Element	Description
LSP Interval Parameters for level 2	
Use level 1 parameter also for level 2	If you want the values you configured for Level 1 to also apply to Level 2, check the Use level 1 parameters also for level 2 check box.
LSP Calculation Interval	Enter the interval of time in seconds between transmission of each LSP. The range is 1-120 seconds. The default is 5. The number should be greater than the expected round-trip delay between any two ASAs on the attached network. The number should be conservative or needless transmission results. Retransmissions occur only when LSPs are dropped. So setting the number to a higher value has little effect on reconvergence. The more neighbors the ASAs have, and the more paths over which LSPs can be flooded, the higher you can make this value.
Initial wait for LSP calculation	Enter the time in milliseconds specifying the initial wait time before the first LSP is generated. The range is 1 to 120,000. The default is 50.
Minimum wait between first and second	Enter the time in milliseconds between the first and second LSP generation. The range is 1 to 120,000. The default is 5000.
Maximum LSP size	In the Maximum LSP size field, enter the number of seconds. The range is 128 to 4352. The default is 1492.
LSP refresh interval	In the LSP refresh interval field, enter the number of seconds at which LSPs are refreshed. The range is 1 to 655535. The default is 900. The refresh interval determines the rate at which the software periodically transmits in LSPs the route topology information that it originates. This is done to keep the database information from becoming too old. Reducing the refresh interval reduces the amount of time that undetected link state database corruption can persist at the cost of increased link utilization. (This is an extremely unlikely event, however, because there are other safeguards against corruption.) Increasing the interval reduces the link utilization caused by the flooding of refreshed packets (although this utilization is very small).
Maximum LSP lifetime	In the Maximum LSP lifetime field, enter the maximum number of seconds that LSPs can remain in a router's database without being refreshed. The range is 1 to 65535. The default is 1200 (20 minutes). You might need to adjust this parameter if you change the LSP refresh interval. LSPs must be periodically refreshed before their lifetimes expire. The value set for LSP refresh interval should be less that the value set for the maximum LSP lifetime; otherwise LSPs will time out before they are refreshed. If you make the LSP lifetime too low compared to the LSP refresh interval, the LSP refresh interval is automatically reduced to prevent the LSPs from timing out.

Summary Address Tab

Use the Add/Edit button to add or edit Summary Addresses.

Field Reference

Table 56-40 ISIS Summary Address Tab

Element	Description
IP address	Enter the IP address of the summary route.
Net Mask	Choose or enter the network mask to apply to the IP address.
Select level	Select the Level 1, Level 2, or Level 1 and 2 radio button depending on which levels you want to receive summary addresses.
Tag	In the Tag field, enter a number for the tag. The range is from 1 to 4294967295.
Metric	In the Metric field. enter the metric that will be applied to the summary route. The range is from 1 to 4294967295. The default value is 10.

Network Entity Title Tab

Use the Add/Edit button to add to edit Network Entity Title.

Field Reference

Table 56-41 ISIS Network Entity Title Tab

Element	Description
Network Entity Title (NET)	Enter a value in the address format 48.0000.1111.2222.00. The total length of NET address must be between 16 and 40 characters.
NET Pool	Click Select to open the NET Pool Object Selector dialog box. You can add and edit NET Pool Objects using this dialog box. For more information about how to add or edit NET Pool Objects, see Add or Edit NET Pool Object Dialog Box, page 6-95. The NET Pool is applicable only for cluster devices in individual mode. Network Entity Title (NET) is not applicable for cluster devices in individual mode.
Maximum allowed NET	Enter a NET value in the range of 3 to 254. The default value is 3.

Interface Tab

Use the Add/Edit button to add or edit ISIS Interfaces.

The ISIS Interface tab has five tabbed panels.

Interface Tab—General Tab

Field Reference

Table 56-42 ISIS Interfaces Tab—General Tab

Element	Description
Interface	Select the interface from available interfaces.
Shutdown ISIS on this interface	Shutdown ISIS on this interface—Lets you disable the IS-IS protocol for this interface without removing the configuration parameters. The IS-IS protocol will not form any adjacencies on this interface and the IP address of this interface will be put into the LSP that is generated by the ASA.
Enable ISIS on this interface	Enables IS-IS protocol on the selected interface.
Enable IPv6 ISIS on this interface	Enables IPv6 IS-IS routing on the selected interface.
Priority for level 1	Lets you set a priority for Level 1. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.
Priority for level 2	Lets you set a priority for Level 2. The priority is used to determine which router on a LAN will be the designated router or Designated Intermediate System (DIS). The priorities are advertised in the hello packets. The router with the highest priority becomes the DIS. The range is 0 to 127. The default is 64.
Tag	Sets a tag on the IP address configured for an interface when this IP prefix is put into an ISIS LSP.
CSNP Interval for level 1	Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 1. This interval only applies for the designated router. The range is from 0 to 65535. The default is 10 seconds.
CSNP Interval for level 2	Sets the Complete Sequence Number PDUs (CSNPs) interval in seconds between transmission of CSNPs on multiaccess networks for Level 2. This interval only applies for the designated router. The range is from 0 to 65535. The default is 10 seconds.
Adjacency filter	Filters the establishment of IS-IS adjacencies.
Match all area addresses	All NSAP addresses must match the filter to accept the adjacency. If not specified (the default), only one address must match the filter for the adjacency to be accepted.

Interface Tab—Authentication Tab

Field Reference

Table 56-43 ISIS Interfaces Tab—Authentication Tab

Element	Description
Level 1 parameters	
Key Type	Select Clear Text or Encrypted.
Key	Enter the key to authenticate IS-IS updates. The range is 0 to 8 characters. If no password is configured with the Key option, no key authentication is performed. **Note** If you selected Key Type as Clear Text, you can enter a maximum of 17 characters in the Key field. If you selected Key Type as Encrypted, you can enter a maximum of 50 characters in the Key field.
Send only	For Send only click the Enable or Disable radio button. Choosing Send only causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.
Mode	Choose the authentication mode by checking the Mode check box and then choosing MD5 or Text from the drop-down list.
Password	Enter a password. **Note** You can select either Mode or enter a password value.
Level 2 parameters	
Key Type	Select Clear Text or Encrypted.
Key	Enter the key to authenticate IS-IS updates. The range is 0 to 8 characters. If no password is configured with the Key option, no key authentication is performed. **Note** If you selected Key Type as Clear Text, you can enter a maximum of 17 characters in the Key field. If you selected Key Type as Encrypted, you can enter a maximum of 50 characters in the Key field.
Send only	For Send only click the Enable or Disable radio button. Choosing Send only causes the system only to insert the password into the SNPs, but not check the password in SNPs that it receives. Use this keyword during a software upgrade to ease the transition. The default is disabled.
Mode	Choose the authentication mode by checking the Mode check box and then choosing MD5 or Text from the drop-down list.
Password	Enter a password. **Note** You can select either Mode or enter a password value.

Interface Tab—Hello Padding Tab

Field Reference

Table 56-44 ISIS Interfaces Tab—Hello Padding Tab

Element	Description
Hello Padding	Enables Hello Padding. IS-IS hellos are padded to the full maximum transmission unit (MTU) size. Padding IS-IS hellos to the full MTU allows for early detection of errors that result from transmission problems with large frames or errors that result from mismatched MTUs on adjacent interfaces.
Minimal holdtime 1 second for level 1	Enables the holdtime (in seconds) that the LSP remains valid for Level 1.
Hello interval for level 1	Specifies the length of time in seconds between hello packets for Level 1. The range is 1 to 65535. The default is 10.
Minimal holdtime 1 second for level 2	Enables the holdtime (in seconds) that the LSP remains valid for Level 2.
Hello interval for level 2	Specifies the length of time in seconds between hello packets for Level 2. The range is 1 to 65535. The default is 10.
Hello multiplier for level 1	Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 1. The advertised hold time in IS-IS hello packets will be set to the hello multiplier times the hello interval. Neighbors will declare an adjacency to this router down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different routers in one area. The range is 3 to 1000. The default is 3.
Hello multiplier for level 2	Specifies the number of IS-IS hello packets a neighbor must miss before the ASA declares the adjacency is down for Level 2. The advertised hold time in IS-IS hello packets will be set to the hello multiplier times the hello interval. Neighbors will declare an adjacency to this router down after not having received any IS-IS hello packets during the advertised hold time. The hold time (and thus the hello multiplier and the hello interval) can be set on a per-interface basis, and can be different between different routers in one area. The range is 3 to 1000. The default is 3.
Configure Circuit Type	Specifies whether the interface is configured for local routing (level 1), area routing (Level 2), or both local and area routing (Level 1-2).

Interface Tab—LSP Settings Tab

Field Reference

Table 56-45 ISIS Interfaces Tab—LSP Settings Tab

Element	Description
Advertise ISIS Prefix	Allows the advertising of IP prefixes of connected networks in the LSP advertisements per IS-IS interface. Disabling this option is an IS-IS mechanism to exclude IP prefixed of connected network from LSP advertisements thereby reducing IS-IS convergence time.
Retransmit Interval	Specifies the amount of time in seconds between retransmission of each IS-IS LSP on a point-to-point link. The number should be greater than the expected round-trip delay between any two routers on the attached network. The range is 0 to 65535. The default is 5.
Retransmit Throttle Interval	Specifies the amount of time in milliseconds between retransmissions on each IS-IS LSP on a point-to-point interface. This option may be useful in very large networks with many LSPs and many interfaces as a way of controlling LSP retransmission traffic. This option controls the rate at which LSPs can be re-sent on the interface. The range is 0 to 65535. The default is 33.
LSP Interval	Specifies the time delay in millisecond between successive IS-IS LSP transmissions. In topologies with a large number of IS-IS neighbors and interfaces, a router may have difficulty with the CPU load imposed by LSP transmission and reception. This option allows the LSP transmission rate (and by implication the reception rate of other systems) to be reduced. The range is 1 to 4294967295. The default is 33.

Interface Tab—Metrics Tab

Field Reference

Table 56-46 ISIS Interfaces Tab—Metrics Tab

Element	Description
Metrics for level 1	
Use maximum metric value	Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations. This is enabled by default.
Default metric	Enter the number for the metric. The range is 1 to 16777214.
Metrics for level 2	
Use maximum metric value	Specifies the metric assigned to the link and used to calculate the cost from each other router via the links in the network to other destinations. This is enabled by default.

Table 56-46 ISIS Interfaces Tab—Metrics Tab

Element	Description
Default metric	Enter the number for the metric. The range is 1 to 16777214.

Passive Interfaces Tab

The Passive Interfaces tab enables you to allow or suppress routing updates on an interface. Only interfaces configured with a name can be suppressed from sending routing updates.

Field Reference

Table 56-47 ISIS Network Entity Title Tab

Element	Description
Passive Interface	Select from the following options: • None—No Interface is selected. • Default—Open the Interfaces Selector dialog to select interfaces that you want to exclude. By default all interfaces are selected. • Specified Interfaces—Open the Interfaces Selector dialog to select interfaces that you want to select and include.

Configuring BFD Routing

The BFD page provides two tabs for configuring BFD (Bidirectional Forwarding Detection) routing on a firewall device. The following topics provide detailed information on configuring BFD.

Navigation Path

- (Device view) Select **Platform > Routing > BFD** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > BFD** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- About BFD, page 56-66
- Create BFD Template, page 56-70
- Add/ Edit BFD Map Dialog Box, page 56-72
- Add/ Edit BFD Interface Dialog Box, page 56-73

About BFD

Bidirectional Forwarding Detection (BFD) is a detection protocol designed to provide fast forwarding path failure detection times for all media types, encapsulations, topologies, and routing protocols. BFD operates in a unicast, point-to-point mode on top of any data protocol being forwarded between two systems. Packets are carried in the payload of the encapsulating protocol appropriate for the media and the network.

BFD provides a consistent failure detection method for network administrators in addition to fast forwarding path failure detection. Because the network administrator can use BFD to detect forwarding path failures at a uniform rate, rather than the variable rates for different routing protocol hello mechanisms, network profiling and planning are easier and reconvergence time is consistent and predictable.

BFD Asynchronous Mode and Echo Function

BFD can operate in asynchronous mode with or without the echo function enabled.

Asynchronous Mode

In asynchronous mode, the systems periodically send BFD control packets to one another, and if a number of those packets in a row are not received by the other system, the session is declared to be down. Pure asynchronous mode (without the Echo function) is useful because it requires half as many packets to achieve a particular detection time as the Echo function requires.

BFD Echo Function

The BFD echo function sends echo packets from the forwarding engine to the directly-connected single-hop BFD neighbor. The echo packets are sent by the forwarding engine and forwarded back along the same path to perform detection. The BFD session at the other end does not participate in the actual forwarding of the echo packets. Because the echo function and the forwarding engine are responsible for the detection process, the number of BFD control packets that are sent out betweenBFD neighbors is reduced. And also because the forwarding engine is testing the forwarding path on the remote neighbor system without involving the remote system, the inter-packet delay variance is improved. This results in quicker failure detection times.

When the echo function is enabled, BFD can use the slow timer to slow down the asynchronous session and reduce the number of BFD control packets that are sent between BFD neighbors, which reduces processing overhead while at the same time delivering faster failure detection.

Note The echo function is not supported for IPv4 multi-hop or IPv6 single-hop BFD neighbors.

You can enable BFD at the interface and routing protocol levels. You must configure BFD on both systems (BFD peers). After you enable BFD on the interfaces and at the router level for the appropriate routing protocols, a BFD session is created, BFD timers are negotiated, and the BFD peers begin to send BFD control packets to each other at the negotiated level.

BFD Session Establishment

The following example shows the ASA and a neighboring router running Border Gateway Protocol (BGP).At the time when both devices come up, there is no BFD session established between them.

Figure 56-1 BFD Session Initiated

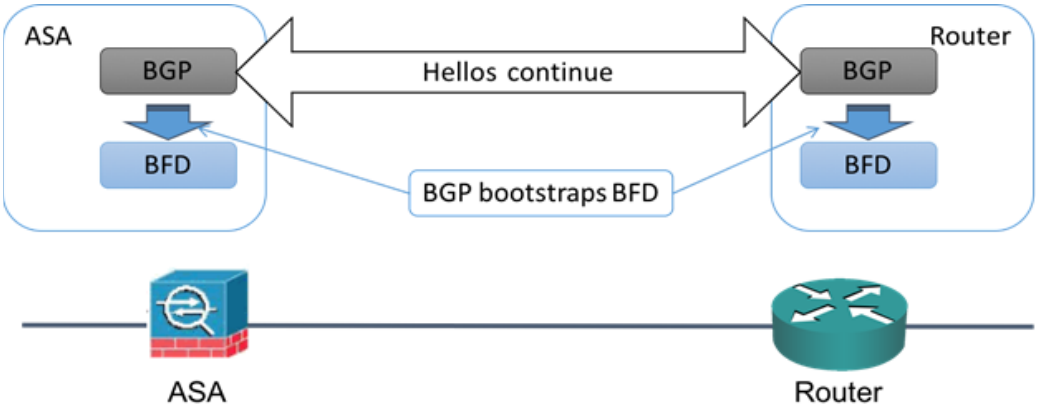

After BGP identifies its BGP neighbor, it bootstraps the BFD process with the IP address of the neighbor. BFD does not discover its peers dynamically. It relies on the configured routing protocols to tell it which IP addresses to use and which peer relationships to form.

The BFD on the router and the BFD on the ASA form a BFD control packet and start sending the packets to each other at a one-second interval until the BFD session is established. The initial control packets from either system are very similar, for example, the Vers, Diag, H, D, P, and F bits are all set to zero, and the State is set to Down. The My Discriminator field is set to a value that is unique on the transmitting device. The Your Discriminator field is set to zero because the BFD session has not yet been established. The TX and RX timers are set to the values found in the configuration of the device.

After the remote BFD device receives a BFD control packet during the session initiation phase, it copies the value of the My Discriminator field into its own Your Discriminator field and the transition from Down state to Init state and then eventually to Up state occurs. Once both systems see their own Discriminators in each other's control packets, the session is officially established.

The following illustration shows the established BFD connection.

Figure 56-2 BFD Session Established

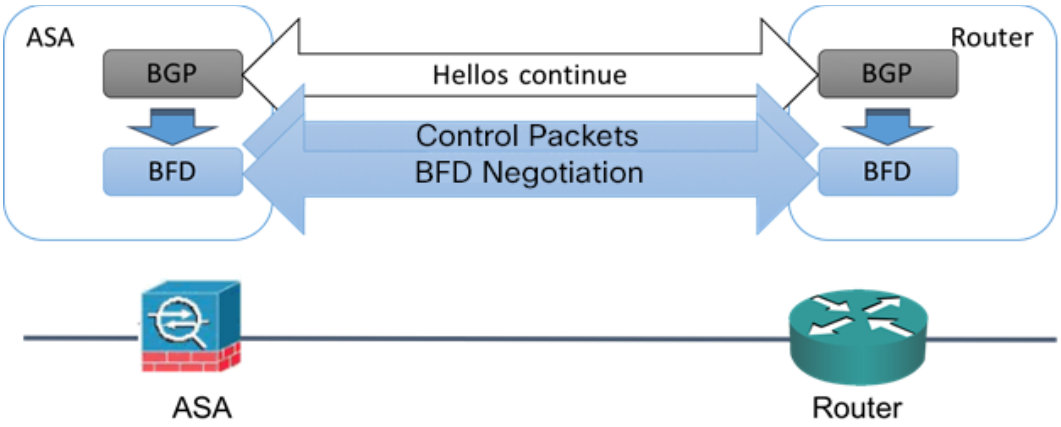

BFD Timer Negotiation

BFD devices must negotiate the BFD timers to control and synchronize the send rate of BFD control packets.

A device needs to ensure the following before it can negotiate a BFD timer:

- That its peer device saw the packet containing the proposed timers of the local device
- That it never sends BFD control packets faster than the peer is configured to receive them
- That the peer never sends BFD control packets faster than the local system is configured to receive them

The setting of the Your Discriminator field and the H bit are sufficient to let the local device that the remote device has seen its packets during the initial timer exchange. After receiving a BFD control packet, each system takes the Required Min RX Interval and compares it to its own Desired Min TX Interval, and then takes the greater (slower) of the two values and uses it as the transmission rate for its BFD packets. The slower of the two systems determines the transmission rate.

When these timers have been negotiated, they can be renegotiated at any time during the session without causing a session reset. The device that changes its timers sets the P bit on all subsequent BFD control packets until it receives a BFD control packet with the F bit set from the remote system. This exchange of bits guards against packets that might otherwise be lost in transit.

Note The setting of the F bit by the remote system does not mean that it accepts the newly proposed timers. It indicates that the remote system has seen the packets in which the timers were changed.

BFD Failure Detection

When the BFD session and timers have been negotiated, the BFD peers send BFD control packets to each other at the negotiated interval. These control packets act as a heartbeat that is very similar to IGP Hello protocol except that the rate is more accelerated.

As long as each BFD peer receives a BFD control packet within the configured detection interval (Required Minimum RX Interval), the BFD session stays up and any routing protocol associated with BFD maintains its adjacencies. If a BFD peer does not receive a control packet within this interval, it

informs any clients participating in that BFD session about the failure. The routing protocol determines the appropriate response to that information. The typical response is to terminate the routing protocol peering session and reconverge and thus bypass a failed peer.

Each time a BFD peer successfully receives a BFD control packet in a BFD session, the detection timer for that session is reset to zero. Thus the failure detection is dependent on received packets and NOT when the receiver last transmitted a packet.

BFD Deployment Scenarios

The following describes how BFD operates in these specific scenarios.

Failover

In a failover scenario, BFD sessions are established and maintained between the active unit and the neighbor unit. Standby units do not maintain any BFD sessions with the neighbors. When a failover happens, the new active unit must initiate session establishment with the neighbor because session information is not synched between active and standby units.

For a graceful restart/NSF scenario, the client (BGP IPv4/IPv6) is responsible for notifying its neighbor about the event. When the neighbor receives the information, it keeps the RIB table until failover is complete. During failover, the BFD and the BGP sessions go down on the device. When the failover is complete, a new BFD session between the neighbors is established when the BGP session comes up.

Spanned EtherChannel and L2 Cluster

In a Spanned EtherChannel cluster scenario, the BFD session is established and maintained between the primary unit and its neighbor. Subordinate units do not maintain any BFD sessions with the neighbors.If a BFD packet is routed to the subordinate unit because of load balancing on the switch, the subordinate unit must forward this packet to the primary unit through the cluster link. When a cluster switchover happens, the new primary unit must initiate session establishment with the neighbor because session information is not synched between primary and subordinate units.

Individual Interface Mode and L3 Cluster

In an individual interface mode cluster scenario, individual units maintain their BFD sessions with their neighbors.

Create BFD Template

This section describes the steps required to create a BFD template policy object. The BFD template specifies a set of BFD interval values. BFD interval values as configured in the BFD template are not specific to a single interface. You can also configure authentication for single-hop and multi-hop sessions. You can enable Echo on single-hop only.

Navigation Path

Select **Manage > Policy Objects**, then select **BFD Template** from the Object Type selector. Right-click inside the work area, then select **New Object** or right-click a row and select **Edit Object**.

Field Reference

Table 56-48 Add/Edit BFD Template

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Config Mode	Specify if there is a single IP hop or multiple IP hops between a BFD source and destination associated with an interface.
Enable Echo	(Optional) Select to enable echo. When enabled, echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection. **Note** This is applicable only for single hop configuration mode.
Interval tab (Optional)	
Interval Type	Specify if you want to define the interval type in microseconds or milliseconds.The default interval type is None.
Transmit and Receive Values Interval Values in Microseconds	This section is enabled, if the Interval type is microseconds. Valid values are between 50000 and 999000 microseconds. Minimum Transmit Value - Enter the minimum transmit interval capability in microseconds. Minimum Receive Value - Enter the minimum receive interval capability microseconds.
Transmit and Receive Values Interval Values in Milliseconds	This section is enabled, if the Interval type is milliseconds. Valid values are between 50 and 999 milliseconds. Minimum Transmit Value - Enter the minimum transmit interval capability in milliseconds. Minimum Receive Value - Enter the minimum receive interval capability milliseconds.
Multiplier Value	Enter the number of consecutive BFD control packets that must be missed before BFD declares that a peer is unavailable.The default value is 3. Valid values are between 3 to 50.
Authentication tab (Optional)	
Authentication Type	Select to configure authentication for the BFD template and specify if you want to use an encrypted password or an unencrypted password,for the authentication.

Table 56-48 Add/Edit BFD Template

Element	Description
Key Value	Enter a BFD password and confirm it. • For **encrypted** BFD templates, the length of the Key value is between 17 and 66 characters. • For **unencrypted** BFD templates of **sha-1 or meticulous-sha-1**authentication type the length of the Key Value must be less than 29 characters. • For **unencrypted** BFD templates of **md5** or **meticulous-md5** authentication type the length of the Key Value must be less than 25 characters.
Key ID	Enter an authentication key ID. This is a shared key ID, that matches the key string.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add/ Edit BFD Map Dialog Box

The Add/ Edit BFD Map dialog box lets you create a BFD map containing destinations that you can associate with a multi-hop template. You must have a multi-hop BFD template already configured. For more information see, Create BFD Template, page 56-70

Navigation Path

You can access the Add/ Edit BFD Map dialog box from the Maps tab on the BFD page. Click the Add Row button to add a new BFD map; select an existing BFD map and click the Edit Row button to edit that map.

Related Topics

- Create BFD Template, page 56-70

Field Reference

Table 56-49 BFD Maps Tab

Element	Description
BFD Template	Select a multi-hop BFD template or add a multi-hop BFD template. For more information see, Create BFD Template, page 56-70
IP version	Select the appropriate address format for the source and destination - IPv4 or IPv6.

Table 56-49 BFD Maps Tab (continued)

Element	Description
IPv4 Destination/Prefix, IPv4 Source/Prefix	Enter the IPv4 address for the destination and source in the appropriate fields in the *x.x.x.x/prefix* format.
IPv6 Destination/ Prefix, IPv6 Source/prefix	Enter the IPv6 address for the destination and source in the appropriate fields in the *x:x:x:x:x:x:x:x/prefix* format.
Slow Timers	This reduces the number of BFD control packets, that are sent between BFD neighbors. This slows down the asynchronous session, reduces the processing overhead and results in faster failure detection. The default value for slow timers is 1000 and valid values are between 1000 - 30000.

Add/ Edit BFD Interface Dialog Box

The Add/ Edit BFD Interfaces dialog box lets you bind a BFD template to an interface, configure the baseline BFD session parameters per interface, and enable echo mode per interface.

Navigation Path

You can access the Add/ Edit BFD Interface dialog box from the Interface tab on the BFD page. Click the Add Row button to add a new BFD interface; select an existing BFD interface and click the Edit Row button to edit that map.

Related Topics

- Create BFD Template, page 56-70

Field Reference

Table 56-50 BFD Interface Tab

Element	Description
Interface	Enter an interface name, select an interface or add an interface role.
BFD Configuration	Select BFD template to select an existing single-hop BFD template or add a single-hop BFD template. Alternately, select BFD interval. For more information see, Create BFD Template, page 56-70
BFD Interval	
Minimum Transmit Value	Enter the minimum transmit interval capability in milliseconds.Valid values are between 50 and 999 milliseconds
Minimum Receive Value	Enter the minimum receive interval capability milliseconds. Valid values are between 50 and 999 milliseconds
Multiplier	Enter the number of consecutive BFD control packets that must be missed before BFD declares that a peer is unavailable.The default value is 3. Valid values are between 3 to 50.
Echo	(Optional) Select to enable echo. When enabled, echo packets are sent by the forwarding engine and forwarded back along the same path in order to perform detection.

Configuring OSPF

The OSPF page provides ten tabbed panels for configuring OSPF (Open Shortest Path First) routing on a firewall device. The following topics provide detailed information about enabling and configuring OSPF:

Note Depending on the device version that you are configuring, some tabs might not be available.

Note Beginning with ASA version 9.2(1), certain OSPF settings have changed. If you configure a shared policy that uses settings specific to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to a device whose version is earlier than 9.2(1). Likewise, if you configure a shared policy that uses settings that no longer apply to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

- About OSPF, page 56-75
- General Tab, page 56-76
- Area Tab, page 56-81
- Range Tab, page 56-84
- Neighbors Tab, page 56-85
- Redistribution Tab, page 56-86
- Virtual Link Tab, page 56-89
- Filtering Tab, page 56-92
- Filter Rule Tab, page 56-94
- Summary Address Tab, page 56-95
- Interface Tab, page 56-96

Navigation Path

- (Device view) Select **Platform > Routing > OSPF** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPF** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

About OSPF

Open Shortest Path First (OSPF) is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. OSPF propagates link-state advertisements (LSAs) rather than routing table updates. Because only LSAs are exchanged, rather than entire routing tables, OSPF networks converge more quickly than RIP networks.

OSPF supports MD5 and clear-text neighbor authentication. Authentication should be used with all routing protocols whenever possible, because route redistribution between OSPF and other protocols (like RIP) can potentially be used by attackers to subvert routing information.

If NAT is used when OSPF is operating on public and private areas, and if address filtering is required, you need to run two OSPF processes—one process for the public areas and one for the private areas.

A router that has interfaces in multiple areas is called an Area Border Router (ABR). A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR).

An ABR uses LSAs to send information about available routes to other OSPF routers. Using ABR type 3 LSA filtering, you can have separate private and public areas with the security appliance acting as an ABR. Type 3 LSAs (inter-area routes) can be filtered from one area to other. This lets you use NAT and OSPF together without advertising private networks.

Note Only type 3 LSAs can be filtered. If you configure the security appliance as an ASBR in a private network, it will send type 5 LSAs describing private networks, which will be broadcast to the entire autonomous system (AS) including public areas.

If NAT is employed but OSPF is only running in public areas, routes to public networks can be redistributed inside the private network, either as default or type 5 AS External LSAs. However, you need to configure static routes for the private networks protected by the security appliance. Also, you should not mix public and private networks on the same security appliance interface.

Related Topics

- Configuring OSPF, page 56-75

General Tab

Use the General panel on the OSPF page to enable up to two OSPF process instances. Each OSPF process has its own associated areas and networks.

Note You cannot enable OSPF if you have RIP enabled.

Navigation Path

You can access the General panel from the OSPF Page; see Configuring OSPF, page 56-75 for more information.

Related Topics

- Area Tab, page 56-81
- Range Tab, page 56-84
- Neighbors Tab, page 56-85
- Redistribution Tab, page 56-86
- Virtual Link Tab, page 56-89
- Filtering Tab, page 56-92
- Summary Address Tab, page 56-95
- Interface Tab, page 56-96

Field Reference

Table 56-51 OSPF General Tab

Element	Description
The General tab provides two identical sections; each is used to enable one OSPF process. The following options are available in each section.	
Enable this OSPF Process	Check this box to enable an OSPF process. You cannot enable an OSPF process if you have RIP enabled on the security appliance. Deselect this option to remove the OSPF process.
OSPF Process ID	Enter a unique numeric identifier for the OSPF process. This process ID is used internally and does not need to match the OSPF process ID on any other OSPF devices. Valid values are from 1 to 65535.
Advanced button	Opens the OSPF Advanced Dialog Box, page 56-77, in which you can configure additional process-related parameters, such as Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings.

OSPF Advanced Dialog Box

Use the OSPF Advanced dialog box to configure settings such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for an OSPF process.

Note Beginning with ASA version 9.2(1), certain OSPF settings have changed. If you configure a shared policy that uses settings specific to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to a device whose version is earlier than 9.2(1). Likewise, if you configure a shared policy that uses settings that no longer apply to ASA 9.2(1)+, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

Navigation Path

You can access the OSPF Advanced dialog box from the General Tab, page 56-76.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-52 OSPF Advanced Dialog Box

Element	Description
OSPF Process	Displays the ID of the OSPF process you are configuring. You cannot change this value in this dialog box.
General Tab	
Router ID	To use a fixed router ID, select **IP Address** and then enter a router ID in IP address format in the Router ID field. To have the router ID automatically generated (the highest-level IP address on the security appliance is used as the router ID), select **Automatic**.

Table 56-52 OSPF Advanced Dialog Box (continued)

Element	Description
Ignore LSA MOSPF	Select this option to suppress transmission of syslog messages when the security appliance receives Type 6 (MOSPF) LSA packets.
RFC 1583 Compatible	Select this option to calculate summary route costs per RFC 1583. Deselect this option to calculate summary route costs per RFC 2328. To minimize the chance of routing loops, all OSPF devices in an OSPF routing domain should have RFC compatibility set identically. This option is selected by default.
Adjacency Changes	These options specify the syslog messages sent when adjacency changes occur. • Log Adjacency Changes – When selected, the security appliance sends a syslog message whenever an OSPF neighbor goes up or down. This option is selected by default. • Log Adjacency Changes Detail – When selected, the security appliance sends a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This option is not selected by default.
Administrative Route Distances	Settings for the administrative route distances, according to the route type. • Inter Area – The administrative distance for all routes from one area to another. Valid values range from 1 to 255; the default value is 110. • Intra Area – The administrative distance for all routes within an area. Valid values range from 1 to 255; the default value is 110. • External – The administrative distance for all routes from other routing domains that are learned through redistribution. Valid values range from 1 to 255; the default value is 110.

Table 56-52 OSPF Advanced Dialog Box (continued)

Element	Description
Timers	Settings used to configure LSA arrival, LSA pacing, and throttling for ASA 9.2(1)+ devices: • LSA Arrival – The minimum delay in milliseconds that must pass between acceptance of the same LSA arriving from neighbors. The range is from 0 to 600,000 milliseconds. The default is 1000 milliseconds. • LSA Flood Pacing – The time in milliseconds at which LSAs in the flooding queue are paced in between updates. The configurable range is from 5 to 100 milliseconds. The default value is 33 milliseconds. • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, checksummed, or aged. Valid values range from 10 to 1800; the default value is 240 seconds. • LSA Retransmission Pacing - The time in milliseconds at which LSAs in the retransmission queue are paced. The configurable range is from 5 to 200 milliseconds. The default value is 66 milliseconds. • LSA Throttle – The delay in milliseconds to generate the first occurrence of the LSA. Valid values range from 0 to 600000 milliseconds. When you enter a value in this field, the Min and Max fields are enabled: – Min – The minimum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. – Max – The maximum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. **Note** For LSA throttling, the first occurrence value must be equal to or less than the minimum value and the minimum value must be equal to or less than the maximum value. • SPF Throttle – The delay to receive a change to the SPF calculation. Valid values range from 1 to 600000 milliseconds. When you enter a value in this field, the Min and Max fields are enabled: – Min – The delay between the first and second SPF calculations. Valid values range from 1 to 600000 milliseconds. – Max – The maximum wait time for SPF calculations. Valid values range from 1 to 600000 milliseconds. **Note** For SPF throttling, the first occurrence value must be equal to or less than the minimum value and the minimum value must be equal to or less than the maximum value. Settings used to configure LSA pacing and SPF calculation timers for device versions earlier than 9.2(1): • SPF Delay – The time between receipt of a topology change and the start of shortest path first (SPF) calculations. Valid values range from 0 to 65535; the default value is 5 seconds. • SPF Hold – The hold time between consecutive SPF calculations. Valid values range from 1 to 65534; the default value is 10 s • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, checksummed, or aged. Valid values range

Table 56-52 OSPF Advanced Dialog Box (continued)

Element	Description
Default Information Originate	Settings used by an ASBR to generate a default external route into an OSPF routing domain. • Enable Default Information Originate – Check this box to enable generation of a default route into the OSPF routing domain; the following options become available: – Always advertise the default route – Check this box to always advertise the default route. – Metric Value – Enter the OSPF metric for the default route. Valid values range from 0 to 16777214; the default value is 1. – Metric Type – Choose the external link type associated with the default route advertised into the OSPF routing domain. The choices are **1** or **2**, indicating a Type 1 or a Type 2 external route. The default value is 2. – Route Map – (Optional) Enter or Select a route map object to apply. The routing process generates the default route if the route map is satisfied. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Non Stop Forwarding Tab **Note** Non Stop Forwarding (NSF) is supported on ASA 9.3(1)+ devices in Spanned Cluster mode or Failover mode only.	
Enable Cisco Non Stop Forwarding Capability	Enables configuration of Cisco nonstop forwarding (NSF) operations.
Enable Cisco Non Stop Forwarding Helper mode	Enables Cisco nonstop forwarding (NSF) helper mode. When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode--the OSPF router process performs nonstop forwarding recovery due to a Route Processor (RP) switchover. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process. If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable Cisco Non Stop Forwarding Helper mode option.
Enable Cisco Non Stop Forwarding	Enables Cisco nonstop forwarding (NSF).

Table 56-52 OSPF Advanced Dialog Box (continued)

Element	Description
Cancel NSF restart when non-NSF-aware neighboring networking devices are detected (Enforce Global)	If neighbors that are not NSF-aware are detected on a network interface during an NSF graceful restart, restart is aborted on that interface only and graceful restart will continue on other interfaces. To cancel restart for the entire OSPF process when neighbors that are not NSF-aware are detected during restart, select the Cancel NSF restart when non-NSF-aware neighboring networking devices are detected (Enforce Global) option. **Note** The NSF graceful restart will also be canceled for the entire process when a neighbor adjacency reset is detected on any interface or when an OSPF interface goes down.
Enable IETF Non Stop Forwarding Capability	Enables configuration of Internet Engineering Task Force (IETF) NSF operations.
Enable IETF Non Stop Forwarding Helper mode	Enables IETF nonstop forwarding (NSF) helper mode. When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode--the OSPF router process performs nonstop forwarding recovery due to a Route Processor (RP) switchover. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process. If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable IETF Non Stop Forwarding Helper mode option.
Enable Strict Link State advertisement checking	Enables strict link-state advertisement (LSA) checking for IETF NSF helper mode.
Enable IETF Non Stop Forwarding	Enables IETF nonstop forwarding (NSF).
Length of graceful restart interval	(Optional) Specifies the length of the graceful restart interval, in seconds. The range is from 1 to 1800. The default is 120. **Note** For a restart interval below 30 seconds, graceful restart will be terminated.

Area Tab

Use the Area tab on the OSPF page to configure OSPF areas and networks.

Navigation Path

You can access the Area tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit Area/Area Networks Dialog Box, page 56-82
- Configuring OSPF, page 56-75

- General Tab, page 56-76
- Range Tab, page 56-84
- Neighbors Tab, page 56-85
- Redistribution Tab, page 56-86
- Virtual Link Tab, page 56-89
- Filtering Tab, page 56-92
- Summary Address Tab, page 56-95
- Interface Tab, page 56-96

Field Reference

Table 56-53 Area Tab

Element	Description
OSPF Process	The OSPF process the area applies to.
Area ID	The area ID.
Area Type	The area type (Normal, Stub, or NSSA).
Networks	The area networks.
Options	The options, if any, set for the area type.
Authentication	The type of authentication set for the area (None, Password, or MD5).
Cost	The default cost for the area.

Add/Edit Area/Area Networks Dialog Box

Use the Add/Edit Area/Area Networks dialog box to define area parameters, the networks contained by the area, and the OSPF process associated with the area.

Navigation Path

You can access the Add/Edit Area/Area Networks dialog box from the Area Tab, page 56-81.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-54 Add/Edit Area/Area Networks Dialog Box

Element	Description
OSPF Process	When adding a new area, choose the OSPF process ID for the OSPF process for which the area is being added. If there is only one OSPF process enabled on the security appliance, that process is selected by default. When editing an existing area, you cannot change the OSPF process ID.

Table 56-54 Add/Edit Area/Area Networks Dialog Box (continued)

Element	Description
Area ID	When adding a new area, enter the area ID. You can specify the area ID as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295. You cannot change the area ID when editing an existing area.
Area Type	
Normal	Choose this option to make the area a standard OSPF area. This option is selected by default when you first create an area.
Stub	Choosing this option makes the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you create a stub area, you can prevent summary LSAs (Type 3 and 4) from being flooded into the area by deselecting the Summary check box.
Summary (allows sending LSAs into the stub area)	When the area being defined is a stub area, deselecting this check box prevents LSAs from being sent into the stub area. This check box is selected by default for stub areas.
NSSA	Choose this option to make the area a not-so-stubby area. NSSAs accept Type 7 LSAs. When you create a NSSA, you can prevent summary LSAs from being flooded into the area by deselecting the Summary check box. You can also disable route redistribution by deselecting the Redistribute check box and enabling Default Information Originate.
Redistribute (imports routes to normal and NSSA areas)	Deselect this check box to prevent routes from being imported into the NSSA. This check box is selected by default.
Summary (allows sending LSAs into the NSSA area)	When the area being defined is a NSSA, deselecting this check box prevents LSAs from being sent into the stub area. This check box is selected by default for NSSAs.
Default Information Originate (generate a Type 7 default)	Select this check box to generate a Type 7 default into the NSSA. This check box is deselected by default.
Metric Value	Specifies the OSPF metric value for the default route. Valid values range from 0 to 16777214. The default value is 1.
Metric Type	The OSPF metric type for the default route. The choices are 1 (Type 1) or 2 (Type 2). The default value is 2.
Network	The IP address and network mask of the network or host to be added to the area. Use 0.0.0.0 with a netmask of 0.0.0.0 to create the default area. You can only use 0.0.0.0 in one area. **Tip** You can click **Select** to select the interfaces from a list of interface objects.

Table 56-54 Add/Edit Area/Area Networks Dialog Box (continued)

Element	Description
Authentication	Contains the settings for OSPF area authentication. • None—Choose this option to disable OSPF area authentication. This is the default setting. • Password—Choose this option to use a clear text password for area authentication. This option is not recommended where security is a concern. • MD5—Choose this option to use MD5 authentication.
Default Cost	Specify a default cost for the area. Valid values range from 0 to 65535 for ASA devices earlier than 9.2(1) and from 0 to 16777214 for ASA 9.2(1)+. The default value is 1.

Range Tab

Use the Range tab to summarize routes between areas.

Navigation Path

You can access the Range tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit Area Range Network Dialog Box, page 56-84

Field Reference

Table 56-55 Range Tab

Element	Description
Process ID	The ID of the OSPF process associated with the route summary.
Area ID	The ID of the area associated with the route summary.
Network	The summary IP address and network mask.
Advertise	Displays “true” if the route summaries are advertised when they match the address/mask pair or “false” if the route summaries are suppressed when they match the address/mask pair.

Add/Edit Area Range Network Dialog Box

Use the Add/Edit Area Range Network dialog box to add a new entry to the Route Summarization table or to change an existing entry.

Navigation Path

You can access the Add/Edit Area Range Network dialog box from the Range Tab, page 56-84.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-56 Add/Edit Area Range Network Dialog Box

Element	Description
OSPF Process	Select the OSPF process to which the route summary applies. You cannot change this value when editing an existing route summary entry.
Area	Select the area ID of the area to which the route summary applies. You cannot change this value when editing an existing route summary entry.
Network	The IP address and mask of the network for the routes being summarized. **Tip** You can click **Select** to select the networks from a list of network objects.
Advertise	Select this check box to set the address range status to "advertise". This causes Type 3 summary LSAs to be generated. Deselect this check box to suppress the Type 3 summary LSA for the specified networks.

Neighbors Tab

Use the Neighbors tab to define static neighbors. You need to define a static neighbor for each point-to-point, non-broadcast interface. You also need to define a static route for each static neighbor in the Neighbors table.

Navigation Path

You can access the Neighbors tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit Static Neighbor Dialog Box, page 56-85

Field Reference

Table 56-57 Neighbors Tab

Element	Description
OSPF Process	The OSPF process associated with the static neighbor.
Neighbor	The IP address of the static neighbor.
Interface	The interface associated with the static neighbor.

Add/Edit Static Neighbor Dialog Box

Use the Add/Edit Static Neighbor dialog box to define a static neighbor or change information for an existing static neighbor. You must define a static neighbor for each point-to-point, non-broadcast interface.

Navigation Path

You can access the Add/Edit Static Neighbor dialog box from the Neighbors Tab, page 56-85.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-58 Add/Edit Static Neighbor Dialog Box

Element	Description
OSPF Process	The OSPF process associated with the static neighbor.
Neighbor	The IP address of the static neighbor. **Tip** You can click **Select** to select the neighbor from a list of host objects.
Interface	The interface associated with the static neighbor. **Tip** You can click **Select** to select the interface from a list of interface objects.

Redistribution Tab

Use the Redistribution tab to define the rules for redistributing routes from one routing domain to another.

Navigation Path

You can access the Redistribution tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Redistribution Dialog Box, page 56-87

Field Reference

Table 56-59 Redistribution Tab

Element	Description
OSPF Process	The OSPF process associated with the route redistribution entry.

Table 56-59 Redistribution Tab (continued)

Element	Description
Route Type	The source protocol the routes are being redistributed from. Valid entries are the following: • BGP—Redistribute routes from the BGP routing process. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS. • EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list. • OSPF—Redistributes routes from another OSPF routing process. • RIP—Redistributes routes from the RIP routing process. • Static—Redistributes static routes to the OSPF routing process.
Match	The conditions used for redistributing routes from one routing protocol to another. These options are not available when redistributing static, connected, RIP, BGP, or EIGRP routes.
Subnets	Displays “true” if subnetted routes are redistributed. Does not display anything if only routes that are not subnetted are redistributed.
Metric Value	The metric that is used for the route. This column is blank for redistribution entries if the default metric is used.
Metric Type	Displays “1” if the metric is a Type 1 external route, “2” if the metric is Type 2 external route.
Tag Value	A 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Route Map	The name of the route map object to apply to the redistribution entry.

Redistribution Dialog Box

Use the Redistribution dialog box to add a redistribution rule or to edit an existing redistribution rule in the Redistribution table.

Navigation Path

You can access the Redistribution dialog box from the Redistribution Tab, page 56-86.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-60 OSPF Redistribution Settings Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the route redistribution entry.
Route Type	Select the source protocol from which the routes are being redistributed. You can choose one of the following options: • BGP—Redistribute routes from the BGP routing process. • Connected—Redistributes connected routes (routes established automatically by virtue of having IP address enabled on the interface) to the OSPF routing process. Connected routes are redistributed as external to the AS. • EIGRP—Redistributes routes from the EIGRP routing process. Choose the autonomous system number of the EIGRP routing process from the list. • OSPF—Redistributes routes from another OSPF routing process. If you choose this protocol, the Match options on this dialog box become visible. These options are not available when redistributing static, connected, RIP, BGP, or EIGRP routes. • RIP—Redistributes routes from the RIP routing process. • Static—Redistributes static routes to the OSPF routing process.
Routing Process ID	The autonomous system (AS) number for the BGP or EIGRP routing process.
Match	If you have chosen OSPF as the Route Type, choose the conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions: • Internal—The route is internal to a specific AS. • External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes. • External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes. • NSSA External 1—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes. • NSSA External 2—Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
Metric Value	The metric value for the routes being redistributed. Valid values range from 1 to 16777214. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified. When redistributing other processes to an OSPF process, the default metric is 20 when no metric value is specified.
Metric Type	Select “1” if the metric is a Type 1 external route, “2” if the metric is a Type 2 external route.

Table 56-60 OSPF Redistribution Settings Dialog Box (continued)

Element	Description
Tag Value	The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Use Subnets	When selected, redistribution of subnetted routes is enabled. Deselect this check box to cause only routes that are not subnetted to be redistributed.
Route Map	Enter or Select a route map object to apply to the redistribution entry. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.

Virtual Link Tab

Use the Virtual Link tab to create virtual links. If you add an area to an OSPF network, and it is not possible to connect the area directly to the backbone area, you need to create a virtual link. A virtual link connects two OSPF devices that have a common area, called the transit area. One of the OSPF devices must be connected to the backbone area.

Navigation Path

You can access the Virtual Link tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit OSPF Virtual Link Configuration Dialog Box, page 56-90

Field Reference

Table 56-61 Virtual Link Tab

Element	Description
OSPF Process	The OSPF process associated with the virtual link.
Area ID	The ID of the transit area.
Peer Router	The IP address of the virtual link neighbor.
Authentication	Displays the type of authentication used by the virtual link: • None—No authentication is used. • Password—Clear text password authentication is used. • MD5—MD5 authentication is used.

Add/Edit OSPF Virtual Link Configuration Dialog Box

Use the Add/Edit OSPF Virtual Link Configuration dialog box to define virtual links or change the properties of existing virtual links.

Navigation Path

You can access the Add/Edit OSPF Virtual Link Configuration dialog box from the Virtual Link Tab, page 56-89.

Related Topics

- Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box, page 56-91
- Configuring OSPF, page 56-75

Field Reference

Table 56-62 Add/Edit OSPF Virtual Link Configuration Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the virtual link.
Area ID	Select the area shared by the neighbor OSPF devices. The selected area cannot be an NSSA or a stub area.
Peer Router	Enter the IP address of the virtual link neighbor.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 10 seconds.
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 5 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value is 1 second.

Table 56-62 Add/Edit OSPF Virtual Link Configuration Dialog Box (continued)

Element	Description
Dead Interval	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535 seconds for ASA devices earlier than 9.2(1) and from 1 to 8192 seconds for ASA 9.2(1)+. The default value of this field is four times the interval set by the Hello Interval field.
Authentication	Contains the OSPF authentication options. • None—Choose this option to disable OSPF authentication. • Area—Choose this option to use the authentication type specified for the area. See Add/Edit Area/Area Networks Dialog Box, page 56-82 for information about configuring area authentication. Area authentication is disabled by default. Therefore, unless you have previously specified an area authentication type, interfaces set to area authentication have authentication disabled until you configure this setting. • Password—Choose this option to use clear text password authentication. This is not recommended where security is a concern. • MD5—Choose this option to use MD5 authentication (recommended).
Authentication Password	Contains the settings for entering the password when password authentication is enabled. • Password—Enter a text string of up to 8 characters. • Confirm—Re-enter the password.
MD5 IDs and Keys	Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID. • MD5 Key ID and MD5 Key Table – MD5 Key ID—A numerical key identifier. Valid values range from 1 to 255. – MD5 Key—An alphanumeric character string of up to 16 bytes.

Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box

Use the Add/Edit OSPF Virtual Link MD5 Configuration dialog box to define MD5 keys for authentication of virtual links.

Navigation Path

You can access the Add/Edit OSPF Virtual Link MD5 Configuration dialog box from the Add/Edit OSPF Virtual Link Configuration Dialog Box, page 56-90.

Related Topics

- Add/Edit OSPF Virtual Link Configuration Dialog Box, page 56-90

- Virtual Link Tab, page 56-89
- Configuring OSPF, page 56-75

Field Reference

Table 56-63 Add/Edit OSPF Virtual Link MD5 Configuration Dialog Box

Element	Description
MD5 Key ID	A numerical key identifier. Valid values range from 1 to 255.
MD5 Key	An alphanumeric character string of up to 16 bytes.
Confirm	Re-enter the MD5 key.

Filtering Tab

Use the Filtering tab to configure the ABR Type 3 LSA filters for each OSPF process. ABR Type 3 LSA filters allow only specified prefixes to be sent from one area to another area and restricts all other prefixes. This type of area filtering can be applied out of a specific OSPF area, into a specific OSPF area, or into and out of the same OSPF areas at the same time.

Benefits

OSPF ABR Type 3 LSA filtering improves your control of route distribution between OSPF areas.

Restrictions

Only type-3 LSAs that originate from an ABR are filtered.

Navigation Path

You can access the Filtering tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit Filtering Dialog Box, page 56-93

Field Reference

Table 56-64 Filtering Tab

Element	Description
OSPF Process	The OSPF process associated with the filter entry.
Area ID	The ID of the area associated with the filter entry.
Prefix List Name	The name of the prefix list.
Filtered Network	The IP address and mask of the network being filtered.
Traffic Direction	Displays “Inbound” if the filter entry applies to LSAs coming in to an OSPF area or “Outbound” if it applies to LSAs going out of an OSPF area.
Sequence #	The sequence number for the filter entry. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.

Table 56-64 Filtering Tab (continued)

Element	Description
Action	Displays “Permit” if LSAs matching the filter are allowed or “Deny” if LSAs matching the filter are denied.
Lower Range	The minimum prefix length to be matched.
Upper Range	The maximum prefix length to be matched.

Add/Edit Filtering Dialog Box

Use the Add/Edit Filtering dialog box to add new filters to the Filter table or to modify an existing filter.

Navigation Path

You can access the Add/Edit Filtering dialog box from the Filtering Tab, page 56-92.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-65 Add/Edit Filtering Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the filter entry.
Area ID	Select the ID of the area associated with the filter entry.
Prefix List Name	Enter or Select the appropriate prefix list object. **Tip** Click **Select** to open the Prefix List Object Selector from which you can select a prefix list object. You can also create new objects from the object Prefix List Object selector. For more information, see Add or Edit Prefix List Object Dialog Box, page 56-146.
Filtered Network	Enter the IP address and mask of the network being filtered.
Traffic Direction	Select the traffic direction to filter. Choose “Inbound” to filter LSAs coming into an OSPF area or “Outbound” to filter LSAs going out of an OSPF area.
Sequence Number	Enter a sequence number for the filter. Valid values range from 1 to 4294967294. When multiple filters apply to an LSA, the filter with the lowest sequence number is used.
Action	Select “Permit” to allow the LSA traffic or “Deny” to block the LSA traffic.
Lower Range	Specify the minimum prefix length to be matched. The value of this setting must be greater than the length of the network mask entered in the Filtered Network field and less than or equal to the value, if present, entered in the Upper Range field.

Table 56-65 Add/Edit Filtering Dialog Box (continued)

Element	Description
Upper Range	Enter the maximum prefix length to be matched. The value of this setting must be greater than or equal to the value, if present, entered in the Lower Range field, or, if the Lower Range field is left blank, greater than the length of the network mask length entered in the Filtered Network field.

Filter Rule Tab

Use the Filter Rule tab to configure rules to filter networks received or transmitted in Open Shortest Path First (OSPF) updates.

Note Filter rules are supported on ASA 9.2(1)+ only.

Navigation Path

You can access the Filter Rule tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Add/Edit Filter Rule Dialog Box, page 56-94

Field Reference

Table 56-66 Filter Rule Tab

Element	Description
Process ID	The OSPF process associated with the filter rule.
ACL	Standard IP access list name. The list defines which networks are to be received and which are to be suppressed in routing updates.
Direction	The direction for the filter rule: • in—The rule filters default route information from incoming routing updates. • out—The rule filters default route information from outgoing routing updates.
Interface	(Optional) The interface to which the filter rule applies.
Routing Process	The routing process: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
Routing Process ID	The identifier for the routing process.

Add/Edit Filter Rule Dialog Box

Use the Add/Edit Filter Rule dialog box to add new filter rules to the Filter Rules table or to modify an existing filter rule.

Note Filter rules are supported on ASA 9.2(1)+ only.

Navigation Path

You can access the Add/Edit Filter Rule dialog box from the Filter Rule Tab, page 56-94.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-67 Add/Edit Filter Rule Dialog Box

Element	Description
OSPF Process	Select the OSPF process associated with the filter rule.
ACL	Select an Access Control List that defines which networks are to be received and which are to be suppressed in routing updates.
Direction	Specify the direction for the filter rule: • in—The rule filters default route information from incoming routing updates. • out—The rule filters default route information from outgoing routing updates.
Interface	(Optional) Specify the interface on which to apply the routing updates. Specifying an interface causes the access list to be applied only to routing updates received on that interface.
Routing Process	Select the routing process for which you want to filter: None, BGP, Connected, EIGRP, OSPF, RIP, or Static.
Routing Process ID	Enter the identifier for the routing process. Applies to BGP, EIGRP, and OSPF routing protocols.

Summary Address Tab

Use the Summary Address tab to configure summary addresses for each OSPF routing process.

Routes learned from other routing protocols can be summarized. The metric used to advertise the summary is the smallest metric of all the more specific routes. Summary routes help reduce the size of the routing table.

Using summary routes for OSPF causes an OSPF ASBR to advertise one external route as an aggregate for all redistributed routes that are covered by the address. Only routes from other routing protocols that are being redistributed into OSPF can be summarized.

Navigation Path

You can access the Summary Address tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Table 56-69 on page 56-96

Field Reference

Table 56-68 Summary Address Tab

Element	Description
Process ID	The OSPF process associated with the summary address.
Network	The IP address and network mask of the summary address.
Tag	A 32-bit decimal value attached to each external route. This value is not used by OSPF itself. It may be used to communicate information between ASBRs.
Advertise	Displays “true” if the summary routes are advertised. Displays “false” if the summary route is not advertised.

Add/Edit Summary Address Dialog Box

Use the Add/Edit Summary Address dialog box to add new entries or to modify existing entries in the Summary Address table.

Navigation Path

You can access the Add/Edit Summary Address dialog box from the Summary Address Tab, page 56-95.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-69 Add/Edit Summary Address Dialog Box

Element	Description
OSPF Process	Choose the OSPF process associated with the summary address. You cannot change this information when editing an existing entry.
Network	The IP address and network mask of the summary address.
Tag	The tag value is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between ASBRs. Valid values range from 0 to 4294967295.
Advertise	When selected, summary routes are advertised. Deselect this check box to suppress routes that fall under the summary address. By default, this check box is selected.

Interface Tab

Use the Interface tab to configure interface-specific OSPF authentication routing properties.

Navigation Path

You can access the Interface tab from the OSPF page. For more information about the OSPF page, see Configuring OSPF, page 56-75.

Related Topics

- Table 56-71 on page 56-98

Field Reference

Table 56-70 Interface Tab

Element	Description
Interface	The name of the interface to which the configuration applies.
Authentication	The type of OSPF authentication enabled on the interface. The authentication type can be one of the following values: • None—OSPF authentication is disabled. • Password—Clear text password authentication is enabled. • MD5—MD5 authentication is enabled. • Area—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled.
Point-to-Point	Displays “true” if the interface is set to non-broadcast (point-to-point). Displays “false” if the interface is set to broadcast.
Cost	The cost of sending a packet through the interface.
Priority	The OSPF priority assigned to the interface.
MTU Ignore	Displays “false” if MTU mismatch detection is enabled. Displays “true” if the MTU mismatch detection is disabled.
Database Filter	Displays “true” if outgoing LSAs are filtered during synchronization and flooding. Displays “false” if filtering is not enabled.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 65535 seconds. The default value is 10 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds. The default value is 1 second.

Table 56-70 Interface Tab (continued)

Element	Description
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgement message. If the router receives no acknowledgement, it resends the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Dead Interval	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. Valid values range from 1 to 65535. The default value of this setting is four times the interval set by the Hello Interval field.
Hello Multiplier (ASA 9.2(1)+ only)	The number of hello packets to be sent per second. Valid values are between 3 and 20.

Add/Edit Interface Dialog Box

Use the Add/Edit Interface dialog box to add OSPF authentication routing properties for an interface or to change an existing entry.

Note Beginning with ASA version 9.2(1), the upper limit for acceptable entries for Hello Interval, Transmit Delay, Retransmit Interval, and Dead Interval has been reduced from 65535 seconds to 8192 seconds. If you configure a shared policy that uses a value over 8192, you will receive a validation error if that policy is assigned to an 9.2(1)+ device.

Navigation Path

You can access the Add/Edit Interface dialog box from the Interface Tab, page 56-96.

Related Topics

- Configuring OSPF, page 56-75

Field Reference

Table 56-71 Add/Edit Interface Dialog Box

Element	Description
Interface	The name of the interface to which the configuration applies.

Table 56-71 Add/Edit Interface Dialog Box (continued)

Element	Description
Authentication	The type of OSPF authentication enabled on the interface. The authentication type can be one of the following values: • No Authentication—OSPF authentication is disabled. • Area Authentication—The authentication type specified for the area is enabled on the interface. Area authentication is the default value for interfaces. However, area authentication is disabled by default. So, unless you previously specified an area authentication type, interfaces showing Area authentication have authentication disabled. • Password Authentication—Clear text password authentication is enabled. • MD5 Authentication—MD5 authentication is enabled.
Authentication Password	Contains the settings for entering the password when password authentication is enabled. • Enter Password—Enter a text string of up to 8 characters. • Confirm—Re-enter the password.
MD5 Key IDs and Keys	Contains the settings for entering the MD5 keys and parameters when MD5 authentication is enabled. All devices on the interface using OSPF authentication must use the same MD5 key and ID. • Key ID—Enter a numerical key identifier. Valid values range from 1 to 255. • Key—An alphanumeric character string of up to 16 bytes. • Confirm—Re-enter the MD5 key. Enter the above values, then click >> to add the key information to the Keys table. Select a key entry and then click << to remove it from the Keys table.
Cost	The cost of sending a packet through the interface.
Priority	The OSPF priority assigned to the interface.
MTU Ignore	When selected, MTU mismatch detection is disabled. Clear this check box to enable MTU mismatch detection.
Database Filter All Out	When selected, outgoing LSAs are filtered during synchronization and flooding. Deselect this check box to disable filtering.
Hello Interval (sec)	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 10 seconds.

Table 56-71 Add/Edit Interface Dialog Box (continued)

Element	Description
Transmit Delay (sec)	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 1 second.
Retransmit Interval (sec)	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a router sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the router receives no acknowledgment, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value is 5 seconds.
Dead Interval (sec)	The interval, in seconds, in which no hello packets are received, causing neighbors to declare a router down. For ASA 9.2(1)+ devices, valid values range from 1 to 8192 seconds. For all other devices, valid values range from 1 to 65535 seconds. The default value of this setting is four times the interval set by the Hello Interval field.
Hello Multiplier (Hello/Sec) (ASA 9.2(1)+ only)	The number of hello packets to be sent per second. Valid values are between 3 and 20. **Note** If you specify a Hello Multiplier, the Hello Interval and Dead Interval values will be ignored. If you entered a value for Hello Interval or Dead Interval, you will be asked to confirm that you want to use the Hello Multiplier instead of the Hello Interval and Dead Interval settings.
Point-to-Point	Displays “true” if the interface is set to non-broadcast (point-to-point). Displays “false” if the interface is set to broadcast.

Configuring OSPFv3

The OSPFv3 page provides two tabbed panels for configuring OSPF (Open Shortest Path First) version 3 routing on a firewall device.

Navigation Path

- (Device view) Select **Platform > Routing > OSPFv3** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPFv3** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

This is the basic procedure for configuring an OSPFv3 process and assigning it to an interface on the OSPFv3 page:

1. On the Process Tab, page 56-103:
   - Specify which of the two processes you are configuring by choosing **Process 1** or **Process 2** from the OSPFv3 Process drop-down list.
   - Check **Enable OSPFv3 Process**.
   - Assign a **Process ID**; any positive integer between 1 and 65535.
   - Use the following features as needed to define the process:
   - **Advanced** button, opening the OSPFv3 Advanced Properties Dialog Box, page 56-104.
   - Area Tab (OSPFv3), page 56-108, for managing area, range, and virtual-link definitions, by means of the Add/Edit Area Dialog Box (OSPFv3), page 56-108, Add/Edit Range Dialog Box (OSPFv3), page 56-110, and Add/Edit Virtual Link Dialog Box (OSPFv3), page 56-111.
   - **Redistribution** panel, for managing route redistribution definitions by means of the Add/Edit Redistribution Dialog Box (OSPFv3), page 56-112.
   - **Summary Prefix** panel, for managing summary-prefix definitions by means of the Add/Edit Summary Prefix Dialog Box (OSPFv3), page 56-113.
2. On the OSPFv3 Interface Tab, page 56-114:
   a. Use the Interface and Neighbor panels to assign the process to a specific interface, using the Add/Edit Interface Dialog Box (OSPFv3), page 56-114 and the Add/Edit Neighbor Dialog Box (OSPFv3), page 56-118.

Related Topics

- About OSPFv3, page 56-101

About OSPFv3

Open Shortest Path First (OSPF) is an interior gateway routing protocol that uses link states rather than distance vectors for path selection. Version 3 is basically OSPFv2 enhanced for IPv6. It is similar to OSPFv2 (see About OSPF, page 56-75), but it is not backward compatible. To use OSPF to route both IPv4 and IPv6v packets, it will be necessary to run both OSPFv2 and OSPFv3 concurrently. They co-exist with each other, but do not interact.

Note OSPFv3 is supported on ASA 9.0+ devices operating in single-context, routed mode only. That is, multiple contexts and transparent mode are not supported.

Think of a link as being an interface on a networking device. A link-state protocol makes its routing decisions based on the states of the links that connect source and destination devices. The state of a link is a description of that interface and its relationship to its neighboring networking devices. This interface information includes the IPv6 prefix/length of the interface, the type of network it is connected to, the devices connected to that network, and so on. This information is propagated in various type of link-state advertisements (LSAs). Because only LSAs are exchanged, rather than entire routing tables, OSPF networks converge more quickly than RIP networks.

The ASA can run two processes of the OSPFv3 protocol simultaneously on different sets of interfaces. You might want to run two processes if you have interfaces that use the same IP addresses (NAT allows these interfaces to co-exist, but OSPFv3 does not allow overlapping addresses). Or you might want to

run one process on the inside interface and another on the outside, redistributing a subset of routes between the two processes. Similarly, you might need to segregate private addresses from public addresses.

You can redistribute routes into an OSPFv3 routing process from another OSPFv3 routing process, a RIP routing process, or from static and connected routes configured on OSPFv3-enabled interfaces.

If NAT is employed but OSPFv3 is only running in public areas, routes to public networks can be redistributed inside the private network, either as default or type 5 AS External LSAs. However, you need to configure static routes for the private networks protected by the security appliance. Also, you should not mix public and private networks on the same security appliance interface.

Differences Between OSPFv2 and OSPFv3

The additional features provided by OSPFv3 over OSPFv2 include the following:

- Use of the IPv6 link-local address for neighbor discovery and other features.
- LSAs expressed as prefix and prefix length.
- Addition of two LSA types.
- Handling of unknown LSA types.
- Protocol processing per link.
- Removal of addressing semantics.
- Addition of flooding scope.
- Support for multiple instances per link.
- Authentication support using the IPSec ESP standard for OSPFv3 routing protocol traffic, as specified by RFC-4552.

Configuration Restrictions

The following are ASA OSPFv3 configuration restrictions:

- To enable OSPFv3 on a specific interface, IPv6 should be enabled on the interface and it must be named.
- Only one OSPFv3 process, with one area and one instance, can be assigned to an interface.
- The Interface neighbor entries take effect only when the OSPFv3 is enabled, and network type should be point-to-point on the specified interface.
- Interface neighbor address must be a link-local address.
- Range value in area Range table should be unique across the area.
- If the area is set to NSSA or stub, the same area cannot be set for virtual-link.
- OSPFv3 redistribution not applicable on the same OSPFv3 process.
- If used in an ASA cluster, OSPFv3 encryption should be disabled.
- The Layer 3 cluster pool is not shared between OSPFv3 and the interface.

Related Topics

- Configuring OSPFv3, page 56-100
- Process Tab, page 56-103
- OSPFv3 Interface Tab, page 56-114

Process Tab

Use the Process tab on the OSPFv3 page to enable and configure up to two OSPFv3 routing processes. Each OSPF process has its own associated areas and networks. For each, at minimum, create an area for OSPFv3, enable an interface for OSPFv3, then redistribute the route into the targeted OSPFv3 routing processes. Note that only single-context mode is supported.

Navigation Path

The Process tab is on the OSPFv3 page.

- (Device view) Select **Platform > Routing > OSPFv3** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > OSPFv3** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101
- Area Tab (OSPFv3), page 56-108
- OSPFv3 Interface Tab, page 56-114

Field Reference

Table 56-72 Process Tab

Element	Description
OSPFv3 Process	Identify which OSPFv3 process you are configuring: choose **Process 1** or **Process 2**. You can enable one or both.
Enable OSPFv3 Process	Check this box to enable the chosen OSPFv3 process. Deselect this option to disable the OSPFv3 process; the process configuration information is retained should you wish to re-enable it later.
Process ID	Enter a unique numeric identifier for this process. The ID can be any positive integer between 1 and 65535. This process ID is used internally and does not need to match the OSPFv3 process ID on any other OSPFv3 devices.
Advanced	Opens the OSPFv3 Advanced Properties Dialog Box, page 56-104, in which you can configure additional process-related parameters, such Router ID, Adjacency Changes, Administrative Route Distances, Timers, Default Information Originate, and Passive Interface settings.
Area	Use the tabs and tables in this panel to manage area, range and virtual-link definitions. See Area Tab (OSPFv3), page 56-108 for more about these definitions.
Redistribution	Use this panel to manage redistribution definitions. See Add/Edit Redistribution Dialog Box (OSPFv3), page 56-112 for more about these definitions.
Summary Prefix	Use this panel to manage summary prefix definitions. See Add/Edit Summary Prefix Dialog Box (OSPFv3), page 56-113 for more about these definitions.

OSPFv3 Advanced Properties Dialog Box

Use the OSPF Advanced dialog box to configure settings such as the Router ID, Adjacency Changes, Administrative Route Distances, Timers, and Default Information Originate settings for an OSPF process.

Navigation Path

You can access the OSPF Advanced dialog box from the Process Tab, page 56-103.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101

Field Reference

Table 56-73 OSPF Advanced Dialog Box

Element	Description
OSPF Process	This read-only field displays the ID of the OSPF process you are configuring.
Router ID	On a single device, choose **Automatic** or **IP Address**. (An address field appears when you choose IP Address.) If you choose Automatic, the highest-level IP address on the security appliance is used as the router ID. To use a fixed router ID, choose IP Address and enter an IPv4 address in the Router ID field. On a device cluster, choose **Automatic** or **Cluster Pool**. (An IPv4 Pool object ID field appears when you choose Cluster Pool.) If you choose Cluster Pool, enter or Select the name of the IPv4 Pool object that is to supply the Router ID address. For more information, see Add or Edit IPv4 Pool Dialog Box, page 6-92.
Ignore LSA MOSPF	Select this option to suppress transmission of syslog messages when the security appliance receives Type 6 (MOSPF) LSA packets.
Adjacency Changes	These options specify the syslog messages sent when adjacency changes occur: • Log Adjacency Changes – When selected, the security appliance sends a syslog message whenever an OSPF neighbor goes up or down. Checking this box enables the Include Details option. • Include Details – When selected, the security appliance sends a syslog message whenever any state change occurs, not just when a neighbor goes up or down. This option is available only when Log Adjacency Changes is checked.

Table 56-73 OSPF Advanced Dialog Box (continued)

Element	Description
Administrative Route Distances	Settings for the administrative route distances, according to the route type. • Inter Area – The administrative distance for all routes from one area to another. Valid values range from 1 to 254; the default value is 110. • Intra Area – The administrative distance for all routes within an area. Valid values range from 1 to 254; the default value is 110. • External – The administrative distance for all routes from other routing domains that are learned through redistribution. Valid values range from 1 to 254; the default value is 110.

Table 56-73 OSPF Advanced Dialog Box (continued)

Element	Description
Timers (in milliseconds)	LSA and SPF throttling provide a dynamic mechanism to slow LSA updates in OSPFv3 during times of network instability, and allow faster OSPFv3 convergence by providing LSA rate limiting. The settings used to configure LSA pacing and SPF calculation timers are: • LSA Arrival – The minimum delay between acceptance of the same LSA arriving from neighbors. Valid values range from 0 to 600000 milliseconds. The default is 1000. • LSA Flood Pacing – The amount of time LSAs in the flooding queue are paced in between updates. Valid values range from 5 to 100 milliseconds. The default value is 33. • LSA Group Pacing – The interval at which LSAs are collected into a group and refreshed, check summed, or aged. Valid values range from 10 to 1800; the default value is 240 milliseconds. • LSA Retransmission Pacing – The length of time at which LSAs in the retransmission queue are paced. Valid values range from 5 to 200 milliseconds. The default value is 66. • LSA Throttle – The delay in milliseconds to generate the first occurrence of the LSA. Valid values range from 0 to 600000 milliseconds. When you enter a value in this field, the min and max fields are enabled: – min – The minimum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. – max – The maximum delay for originating the same LSA. Valid values range from 1 to 600000 milliseconds. • SPF Throttle – The delay to receive a change to the SPF calculation. Valid values range from 1 to 600000 milliseconds. When you enter a value in this field, the min and max fields are enabled: – min – The delay between the first and second SPF calculations. Valid values range from 1 to 600000 milliseconds. – max – The maximum wait time for SPF calculations. Valid values range from 1 to 600000 milliseconds. **Note** For LSA throttling, if the minimum or maximum time is less than the first occurrence value, then OSPFv3 automatically corrects to the first occurrence value. Similarly, if the maximum delay specified is less than the minimum delay, then OSPFv3 automatically corrects to the minimum delay value.

Table 56-73 OSPF Advanced Dialog Box (continued)

Element	Description
Default Information Originate	Settings used by an ASBR to generate a default external route into an OSPFv3 routing domain: • Enable Default Information Originate – Check this box to enable generation of a default route into the OSPFv3 routing domain; the following options become available: – Always advertise the default route – Check this box to always advertise the default route. – Metric Value – The OSPFv3 metric used to generate the default route. Valid values range from 0 to 16777214. – Metric Type – The external link type associated with the default route advertised into the OSPFv3 routing domain. Choose **1** or **2**, indicating a Type 1 or a Type 2 external route. The default value is 1. – Route Map – (Optional) Enter or Select the name of a route map object to apply. The routing process generates the default route if the route map is satisfied. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.
Passive Interface	Passive routing helps control the advertisement of OSPFv3 routing information, and disables sending and receiving OSPFv3 routing updates on an interface. Enter or Select one or more interfaces, or interface objects, to enable passive OSPFv3 routing on those interfaces. IPv4 and IPv6 addresses are supported.
Non Stop Forwarding Tab **Note** Non Stop Forwarding (NSF) is supported on ASA 9.3(1)+ only.	
Enable graceful-restart helper	Enables graceful restart helper mode. When an ASA has NSF enabled, it is said to be NSF-capable and will operate in graceful restart mode. By default, the neighboring ASAs of the NSF-capable ASA will be NSF-aware and will operate in NSF helper mode. When the NSF-capable ASA is performing graceful restart, the helper ASAs assist in the nonstop forwarding recovery process. If you do not want the ASA to help the restarting neighbor with nonstop forwarding recovery, clear the Enable graceful-restart helper option.

Table 56-73 OSPF Advanced Dialog Box (continued)

Element	Description
Enable Link State Advertisement	Enables strict link-state advertisement (LSA) checking. **Note** When enabled, it indicates that the helper router will terminate the process of restarting the router if it detects that there is a change to a LSA that would be flooded to the restarting router, or if there is a changed LSA on the retransmission list of the restarting router when the graceful restart process is initiated.
Enable graceful-restart (Use when Spanned Cluster or Failover configured)	Enables graceful restart on the ASA.
Length of graceful restart interval	(Optional) Specifies the length of the graceful restart interval, in seconds. The range is from 1 to 1800. The default is 120. **Note** For a restart interval below 30 seconds, graceful restart will be terminated.

Area Tab (OSPFv3)

Use the Area panel on the Process Tab, page 56-103 of the OSPFv3 page to configure OSPFv3 areas, ranges and virtual links. The Area panel consists of three definition tables—Area, Range, and Virtual Link:

- Refer to Add/Edit Area Dialog Box (OSPFv3), page 56-108 for information about adding and editing Area table entries.
- Refer to Add/Edit Range Dialog Box (OSPFv3), page 56-110 for information about adding and editing Range table entries.
- Refer to Add/Edit Virtual Link Dialog Box (OSPFv3), page 56-111 for information about adding and editing Virtual Link table entries.

Refer to Using Tables, page 1-48 for basic information about working with Security Manager tables.

Navigation Path

You can access the Area tab from the Process Tab, page 56-103 of the OSPFv3 page. For more information about the OSPFv3 page, see Configuring OSPFv3, page 56-100.

Related Topics

- About OSPFv3, page 56-101
- OSPFv3 Interface Tab, page 56-114

Add/Edit Area Dialog Box (OSPFv3)

Use the Add/Edit Area dialog box to define parameters for the area.

Navigation Path

You can access the Add/Edit Area dialog box from the Area Tab (OSPFv3), page 56-108.

Related Topics

- Configuring OSPFv3, page 56-100

- About OSPFv3, page 56-101
- Process Tab, page 56-103

Field Reference

Table 56-74 Add/Edit Area Dialog Box

Element	Description
Area ID	Enter an identifier for the area as either a decimal number or an IP address. Valid decimal values range from 0 to 4294967295.
Cost	The cost of sending a packet on an interface. Valid values are 0 to 65535. Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The ASA calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.
Type	Define the area type by choosing one of the following: • Normal – Make the area a standard OSPFv3 area. This option is selected by default when you first create an area. • NSSA – Make the area a "not-so-stubby area." NSSAs accept Type 7 LSAs. When you choose this option, the Default Information Originate options are enabled. When you create a NSSA, you can prevent summary LSAs from being flooded into the area by deselecting *Allow sending summary LSA into this area*. You can also disable route redistribution by deselecting *Redistribute*, and enabling *Default information originate*. • Stub – Make the area a stub area. Stub areas do not have any routers or areas beyond it. Stub areas prevent AS External LSAs (Type 5 LSAs) from being flooded into the stub area. When you choose this option, *Allow sending summary LSA into this area* is enabled. When you create a stub area, you can prevent summary LSAs (Type 3 and 4) from being flooded into the area by deselecting *Allow sending summary LSA into this area*.
Default Information Originate	
These options are enabled when you choose NSSA as the area Type. The first option is enabled when you choose Stub as the area Type.	
Allow sending summary LSA into this area	Select to allow flooding of summary LSAs into the area.
Redistribute (imports routes to normal and NSSA areas)	Select to allow route redistribution.

Table 56-74 Add/Edit Area Dialog Box (continued)

Element	Description
Default information originate	Check this box to generate a Type 7 default into the NSSA. Selecting this option enables the following metric options: • Metric – The OSPF metric value for the default route. Valid values range from 1 to 16777214. The default is 1. • Metric Type – The OSPF metric type for the default route. Choose 1 (Type 1) or 2 (Type 2). The default is 1.

Add/Edit Range Dialog Box (OSPFv3)

Use the Add/Edit Area Range Network dialog box to add a new range to the area selected in the Area table, or to change an existing entry.

Navigation Path

You can access the Add/Edit Range dialog box from the Range panel under the Area Tab (OSPFv3), page 56-108.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101
- Process Tab, page 56-103

Field Reference

Table 56-75 Add/Edit Range Dialog Box

Element	Description
Area ID	This read-only entry is the ID of the area to which this range applies.
IPv6 Prefix/Length	The IPv6 address(es) for the routes being summarized. **Tip** You can click **Select** to select the networks from a list of network objects.
Cost	The cost for the summary route, which is used during OSPF SPF calculations to determine the shortest paths to the destination. Valid values are 0 to 16777215. Routing decisions are based on cost, which is an indication of the overhead required to send packets across a certain interface. The ASA calculates the cost of an interface based on link bandwidth rather than the number of hops to the destination. The cost can be configured to specify preferred paths.
Advertise	Select this option to set the address range status to advertise. This causes Type 3 summary LSAs to be generated (this is the default). Deselect this optionn to suppress the Type 3 summary LSAs for the specified networks.

Add/Edit Virtual Link Dialog Box (OSPFv3)

Use the Add/Edit Virtual Link dialog box to define virtual links for the area selected in the Area table, or change the properties of existing virtual links.

Navigation Path

You can access the Add/Edit Virtual Link dialog box from the Virtual Link panel under the Area Tab (OSPFv3), page 56-108.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101
- Process Tab, page 56-103

Field Reference

Table 56-76 Add/Edit Virtual Link Dialog Box

Element	Description
Area ID	This read-only entry is the ID of the area to which this virtual link applies.
Peer Router ID	Enter the IP address of the virtual link neighbor. **Tip** You can click **Select** to select from a list of network objects.
TTL Security	The time-to-live (TTL) security hop count on a virtual link. The hop count value can range from 1 to 254.
Dead Interval	The interval, in seconds, if no hello packets are received, neighbors declare the device down. Valid values range from 1 to 8192. The default value of this field is four times the Hello Interval.
Hello Interval	The interval, in seconds, between hello packets sent on an interface. The smaller the hello interval, the faster topological changes are detected but the more traffic is sent on the interface. This value must be the same for all routers and access servers on a specific interface. Valid values range from 1 to 8192 seconds. The default value is 10 seconds.
Transmit Interval	The time, in seconds, between LSA retransmissions for adjacencies belonging to the interface. When a device sends an LSA to its neighbor, it keeps the LSA until it receives the acknowledgment message. If the device does not receive an acknowledgment, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 8192 seconds. The default value is 5 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 8192 seconds. The default value is 1 second.

Add/Edit Redistribution Dialog Box (OSPFv3)

Use the Add/Edit Redistribution dialog box to add a redistribution rule to this process, or to edit an existing redistribution rule.

Navigation Path

You can access the Redistribution dialog box from the Redistribution panel under the Process Tab, page 56-103.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101

Field Reference

Table 56-77 Add/Edit Redistribution Dialog Box

Element	Description
Source Protocol	Choose the source protocol for route redistribution: • Connected – Redistributes connected routes (routes established automatically by virtue of having an IP address enabled on the interface) to the OSPFv3 routing process. Connected routes are redistributed as external to the autonomous system. • OSPF – Redistributes routes from another OSPF routing process. The Routing PID and the Match options are enabled when you choose this option. • Static – Redistributes static routes to the OSPFv3 routing process.
Metric	The metric value for the routes being redistributed. Valid values range from 1 to 16777214; the default is 20. When redistributing from one OSPF process to another OSPF process on the same device, the metric will be carried through from one process to the other if no metric value is specified.
Metric Type	The metric type is the external link type associated with the default route that is advertised into the OSPFv3 routing domain. Choose None, 1, or 2, where None indicates there is no default route, 1 indicates the metric is a Type 1 external route, and 2 is a Type 2 external route.
Tag (optional)	The tag is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between other border devices. Valid values range from 0 to 4294967295.
Route Map	Enter or Select the name of the route map object to apply to the redistribution entry. **Tip** Click **Select** to open the Route Map Object Selector from which you can select a route map object. You can also create new route map objects from the Route Map Object Selector. For more information, see Understanding Route Map Objects, page 56-132.

Table 56-77 Add/Edit Redistribution Dialog Box (continued)

Element	Description
Routing PID	The ID of the process to which redistribution is directed. (The Process ID is defined on the Process Tab, page 56-103.) This option is enabled only when OSPF is chosen as the Source Protocol.
Include Connected	Check this box to include connected routes in the redistribution.
Match The conditions used for redistributing routes from one routing protocol to another. The routes must match the selected condition to be redistributed. You can choose one or more of the following match conditions. These options are enabled only when OSPF is chosen as the Source Protocol.	
Internal	The route is internal to a specific autonomous system.
External 1	Routes that are external to the autonomous system, but are imported into OSPF as Type 1 external routes.
External 2	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 external routes.
NSSA External 1	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.
NSSA External 2	Routes that are external to the autonomous system, but are imported into OSPF as Type 2 NSSA routes.

Add/Edit Summary Prefix Dialog Box (OSPFv3)

Use the Add/Edit Summary Prefix dialog box to add new route-summarization entries to the selected process, or to modify existing entries.

Navigation Path

You can access the Add/Edit Summary Prefix dialog box from the Summary Prefix panel under the Process Tab, page 56-103.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101

Field Reference

Table 56-78 Add/Edit Summary Prefix Dialog Box

Element	Description
Process ID	This read-only value identifies the process to which this rule applies.
IPv6 Prefix/Length	Enter an IPv6 prefix/length for external route summarization. **Tip** You can click **Select** to select from a list of network objects.

Table 56-78 Add/Edit Summary Prefix Dialog Box (continued)

Element	Description
Advertise	When selected, summary routes that match the specified prefix and mask pair are advertised. When deselected, routes that match the specified prefix and mask pair are suppressed. By default, this check box is selected.
Tag (optional)	The tag is a 32-bit decimal value attached to each external route. This is not used by OSPF itself. It may be used to communicate information between border devices. Valid values range from 0 to 4294967295. This field is enabled when you check Advertise.

OSPFv3 Interface Tab

Use the Interface panel to configure interface-and neighbor-specific OSPFv3 routing properties. The Interface panel consists of two definition tables, Interface and Neighbor:

- Refer to Add/Edit Interface Dialog Box (OSPFv3), page 56-114 for information about adding and editing Interface table entries.
- Refer to Add/Edit Neighbor Dialog Box (OSPFv3), page 56-118 for information about adding and editing Neighbor table entries.

Refer to Using Tables, page 1-48 for basic information about working with Security Manager tables.

Navigation Path

Click the Interface tab on the OSPFv3 page to display this panel. For more information about the OSPFv3 page, see Configuring OSPFv3, page 56-100.

Related Topics

- About OSPFv3, page 56-101
- Process Tab, page 56-103

Add/Edit Interface Dialog Box (OSPFv3)

Use the Add/Edit Interface dialog box to define OSPFv3 routing properties for an individual interface, or to change an existing entry.

Navigation Path

You can access the Add/Edit Interface dialog box from the Interface panel under the OSPFv3 Interface Tab, page 56-114.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101
- Process Tab, page 56-103

Field Reference

Table 56-79 Add/Edit Interface Dialog Box

Element	Description
Interface	The name of the interface to which this routing configuration applies. **Tip** You can click **Select** to select from a list of interface objects.
Enable OSPFv3 on this interface	Check this box to enable OSPFv3 on the specified interface, and activate the following fields: • Process ID – Choose the process to apply to this interface; defined on the OSPFv3 Process Tab, page 56-103. • Area ID – Identify the area to be assigned; areas are also defined on the OSPFv3 Process Tab, page 56-103. • Instance ID – (Optional) Specify an ID for this process instance. Valid values for this setting range from 0 to 255. This feature lets you have multiple OSPFv3 processes on a single link. Received packets with other instance IDs are then ignored by this process.
Properties	
Filter outgoing link-state advertisements	Check this box to filter outgoing LSAs. All outgoing LSAs are flooded to the interface by default.
Disable MTU mismatch detection	Check this box to disable the OSPFv3 MTU mismatch detection when database description (DBD) packets are received.
Flood Reduction	Check this box to suppress unnecessary flooding of LSAs in stable topologies.
Point-to-point Network	Check this box to define this as a link to a point-to-point network; that is, a network between two routing devices. All neighbors on a point-to-point network establish adjacency and there is no designated router. This option is unavailable when the Broadcast option is selected.
Broadcast	Check this box to define this as a link to a network with multiple routing devices. Such networks establish a designated router (DR), as well as a backup designated router (BDR), that controls LSA flooding on the network. This option is unavailable when the Point-to-point Network option is selected.
Cost	The cost of sending a packet through the interface. Link cost is an arbitrary number used in shortest path first calculations. If you do not assign a value, the configured reference bandwidth divided by the interface port speed is used. (The default reference bandwidth is 40 Gb/sec.)

Table 56-79 Add/Edit Interface Dialog Box (continued)

Element	Description
Priority	Assign an OSPFv3 priority to this interface. Valid values for this setting range from 0 to 255. Entering 0 for this setting makes the device ineligible to become the designated router or backup designated router. This setting does not apply to interfaces that are configured as point-to-point, non-broadcast interfaces. When two routing devices connect to a network, both attempt to become the designated router. The device with the higher priority becomes the designated router. If there is a tie, the router with the higher router ID becomes the designated router.
Dead Interval	If no hello packets are received from a neighbor within this interval, that device is designated as inactive. Valid values range from 1 to 65535. The default value for this setting is four times the hello interval.
Poll Interval	If a neighboring device is inactive, it may be necessary to continue sending hello packets to that neighbor. The hello packets are sent at this reduced interval, which should be larger than the hello interval.
Retransmit Interval	The time, in seconds, between LSA retransmissions for adjacent neighbors. When a router sends an LSA to a neighbor, it keeps the LSA until it receives an acknowledgment. If an acknowledgment is not received within this interval, it will resend the LSA. Be conservative when setting this value, or needless retransmission can result. The value should be larger for serial lines and virtual links. Valid values range from 1 to 65535 seconds.
Transmit Delay	The estimated time, in seconds, required to send an LSA packet on the interface. LSAs in the update packet have their ages increased by the amount specified by this field before transmission. If the delay is not added before transmission over a link, the time in which the LSA propagates over the link is not considered. The value assigned should take into account the transmission and propagation delays for the interface. This setting has more significance on very low-speed links. Valid values range from 1 to 65535 seconds.
Authentication	
Type	The type of authentication enabled on this interface. Choose one of the following: • Area – OSPFv3 does not provide “built-in” authentication, instead relying on IPv6/IPSec protocols. Choose this option to use those protocols to authenticate OSPFv3 traffic on all interfaces in the area; this means all routing devices in the area must use this option. This is the default. • Interface – Choose this option to secure this interface and protect OSPFv3 virtual links. The additional parameters in this section are enabled when you choose this option. • None – OSPFv3 authentication is disabled.

Table 56-79 Add/Edit Interface Dialog Box (continued)

Element	Description
Security Parameter Index	Enter an IPSec identification tag used to distinguish this particular OSPFv3 interface; used in conjunction with the specified authentication and encryption rules. Valid values range from 256 to 4294967295.
Authentication Algorithm	Choose the type of authentication algorithm to use: • md5 – Message Digest 5; produces a 128-bit hash value. • sha1 – Secure Hash Algorithm version 1; produces a 160-bit hash value.
Authentication Key	Enter an authentication key. The length of the key entered depends on the type of authentication chosen as the Authentication Algorithm, and whether the key is to be encrypted (when you check the Encrypt Authentication Key box): • md5 – 32 characters. • md5 (encrypted) – 66 characters. • sha1 – 40 characters. • sha1 (encrypted) – 82 characters.
Encrypt Authentication Key	Check this box to require encryption of the specified Authentication Key for transmission.
Include Encryption	Check this box to require encryption of OSPFv3 packets. The following options are enabled.
Encryption Algorithm	Choose the type of encryption to use: • 3des – Triple DES; the Data Encryption Standard cipher algorithm is applied three times to each packet. • aes-cbc – Encryption is based on the Advanced Encryption Standard with Cipher Block Chaining, to produce a key of the size chosen with the Key Type parameter. The Key Type list is enabled only when you choose this encryption option. Choose one of these options: – 128 – For 128-bit keys. – 192 – For 192-bit keys. – 256 – For 256-bit keys. • des – Encryption is based on the Data Encryption Standard, using 56-bit keys.

Table 56-79 Add/Edit Interface Dialog Box (continued)

Element	Description
Encryption Key	Enter an encryption key. The length of the key entered depends on the type of encryption chosen as the Encryption Algorithm, and whether the key is to be encrypted (when you check the Encrypt Key box): • 3des – 48 characters (192 bits). • 3des (encrypted) – 98 characters (192 bits). • aes-cbc/128 – 32 characters (128 bits). • aes-cbc/128 (encrypted) – 66 characters (128 bits). • aes-cbc/192 – 48 characters (192 bits). • aes-cbc/192 (encrypted) – 98 characters (192 bits). • aes-cbc/256 – 64 characters (256 bits). • aes-cbc/256 (encrypted) – 130 characters (256 bits). • des – 16 characters (64 bits). • des (encrypted) – 34 characters (64 bits).
Encrypt Key	Check this box to require encryption of the specified Encryption Key for transmission.

Add/Edit Neighbor Dialog Box (OSPFv3)

You must define a static neighbor for each point-to-point, non-broadcast interface. This feature lets you broadcast OSPFv3 advertisements across an existing VPN connection without having to encapsulate the advertisements in a GRE tunnel. Note the following restrictions:

- You cannot define the same static neighbor for two different OSPFv3 processes.
- You must define a static route for each static neighbor.

Use the Add/Edit Neighbor dialog box to define a static neighbor for the interface selected in the Interface table, or to change information for an existing static neighbor.

Navigation Path

You can access the Add/Edit Neighbor dialog box from the Neighbor panel under the OSPFv3 Interface Tab, page 56-114.

Related Topics

- Configuring OSPFv3, page 56-100
- About OSPFv3, page 56-101
- Process Tab, page 56-103

Field Reference

Table 56-80 Add/Edit Neighbor Dialog Box

Element	Description
Interface	The interface associated with this neighbor definition (read-only).

Table 56-80 Add/Edit Neighbor Dialog Box (continued)

Element	Description
Link-local Address	Enter the IPv6 address of the static neighbor.
Cost and Database Filter	Check this box to enable filtering of the outgoing LSAs on the interface during synchronization and flooding. The following options are enabled: • Cost – Use this field to assign an arbitrary cost to this neighbor. If a value is not assigned, the cost of the interface is used (this value is based on the port speed of the interface, and is calculated as reference bandwith divided by interface speed). Valid values range from 1 to 65535. • Filter outgoing link-state advertisements – Check this box to disable forwarding of outgoing LSAs to this neighbor. **Note** The Cost and Database Filter options and the Poll-Interval options are mutually exclusive.
Poll-Interval	Check this box to enable the following options: • Poll Interval – Time interval in seconds between transmission of hello packets to a "dead" neighbor. The default is 120. If a neighboring device becomes inactive (hello packets have not been received for the dead interval period), it may be necessary to continue sending hello packets to the dead neighbor at a reduced rate. Thus this value should be larger than the interface hello interval. • Priority – The router priority value of this neighbor. The default is 0; valid values range from 1 to 255. The priority value helps determine the designated router for an OSPFv3 link. A value of zero means the device is ineligible to become the designated router, or backup designated router. **Note** The Poll-Interval options and the Cost and Database Filter options are mutually exclusive. Also, these values do not apply to point-to-multipoint interfaces.

Configuring RIP

Routing Information Protocol (RIP) is a dynamic routing protocol, or more precisely, an interior gateway protocol that is based on distance vectors. RIP uses hop count as the metric for path selection. When RIP is enabled on an interface, the interface exchanges RIP broadcast packets with neighboring devices to dynamically learn about and advertise routes. These RIP packets contain information about the destination networks that the gateways can reach, and the number of gateways that a packet must travel through to reach those destinations.

Cisco Security Manager supports both RIP version 1 and RIP version 2. Version 1 does not send the subnet mask with the routing update; RIP version 2 sends the subnet mask with the routing update, and supports variable-length subnet masks. Additionally, RIP version 2 supports neighbor authentication when routing updates are exchanged. This authentication ensures that the security appliance receives reliable routing information from a trusted source.

Note You cannot enable RIP if you have OSPF processes running.

Limitations

RIP has the following limitations:

- Cisco Security Manager cannot pass RIP updates between interfaces.
- RIP Version 1 does not support variable-length subnet masks.
- RIP has a maximum hop count of 15. A route with a hop count greater than 15 is considered unreachable.
- RIP convergence is relatively slow compared to other routing protocols.

RIP Version 2 Notes

The following information applies to RIP Version 2 only:

- If using neighbor authentication, the authentication key and key ID must be the same on all neighbor devices that provide RIP version 2 updates to the interface.
- With RIP version 2, the security appliance transmits and receives default route updates using the multicast address 224.0.0.9. In passive mode, it receives route updates at that address.
- When RIP version 2 is configured on an interface, the multicast address 224.0.0.9 is registered on that interface. When a RIP version 2 configuration is removed from an interface, that multicast address is unregistered.

Using Security Manager to Configure RIP on Security Appliances

Use the RIP page to enable the Routing Information Protocol on an interface. The settings and features available when configuring RIP depend on the type of device and OS version that you are configuring:

- To configure RIP on a PIX Firewall or ASA running an OS version earlier than 7.2, or on any FWSM, see RIP Page for PIX/ASA 6.3–7.1 and FWSM, page 56-120.
- To configure RIP on a PIX Firewall or ASA running OS version 7.2 or later, see RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Related Topics

- Configuring Static Routes, page 56-128
- Configuring OSPF, page 56-75
- Configuring No Proxy ARP, page 56-1
- Configuring Routing Information Protocol – a chapter from the "Cisco IOS IP Configuration Guide, Release 12.2," providing additional detailed information about RIP

RIP Page for PIX/ASA 6.3–7.1 and FWSM

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use this RIP page to enable the Routing Information Protocol (RIP) on an interface in any FWSM, or in a PIX/ASA running a pre-7.2 version operating system.

The RIP table on this page lists all interfaces on which RIP is currently defined. Use the Add RIP Configuration and Edit RIP Configuration dialog boxes to create and maintain these entries. See RIP Page for PIX/ASA 6.3–7.1 and FWSM, page 56-120 for more information.

Navigation Path

- (Device view) Select **Platform > Routing > RIP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > RIP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

  When creating a shared RIP policy, you must choose a Version in the Create a Policy dialog box, as follows:

  - **PIX/ASA 6.3-7.1 and FWSM**
  - **PIX/ASA 7.2 and Later**

  When assigning a shared RIP policy, be sure to assign the appropriate RIP policy for the device. For example, you cannot assign a PIX/ASA 7.2+ RIP policy to an FWSM.

Related Topics

- Configuring Static Routes, page 56-128
- Configuring OSPF, page 56-75
- Configuring No Proxy ARP, page 56-1
- RIP Page for PIX/ASA 7.2 and Later, page 56-122
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Table Columns and Column Heading Features, page 1-49

Add/Edit RIP Configuration (PIX/ASA 6.3–7.1 and FWSM) Dialog Boxes

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use the Add RIP Configuration and Edit RIP Configuration dialog boxes to add a RIP configuration to the security appliance, or to make changes to an existing RIP configuration. By adding a RIP configuration, you enable RIP on the specified interface. Except for their titles, the two dialog boxes are identical.

Navigation Path

You can access the Add and Edit RIP Configuration dialog boxes from the RIP Page for PIX/ASA 6.3–7.1 and FWSM, page 56-120.

Field Reference

Table 56-81 Add/Edit RIP Configuration (PIX/ASA 6.3-7.1 and FWSM) Dialog Boxes

Element	Description
Interface	Enter or Select the interface for the RIP configuration. You cannot configure two different RIP configurations on the same interface.

Table 56-81 Add/Edit RIP Configuration (PIX/ASA 6.3-7.1 and FWSM) Dialog Boxes (continued)

Element	Description
Mode	Select the interface behavior regarding RIP updates: • **Send default routes** – The interface will transmit RIP routing updates only. • **Receive routes** – The interface will listen for RIP routing broadcasts and use that information to populate its routing table, but it will not send RIP routing updates. • **Send default routes and receive routes** – The interface will send and receive RIP routing updates.
Version	Select the RIP version to enable on the interface: • **RIP Version 1** – Enables RIP Version 1 on the interface. • **RIP Version 2** – Enables RIP Version 2 on the interface. Configuring RIP Version 2 registers the multicast address 224.0.0.9 on the interface.
Version 2 Authentication	These options let you enable and select the type of authentication used with RIP Version 2. • **Enable Authentication** – This option is available when you select RIP Version 2 above. When this box is checked, RIP neighbor authentication is enabled and the following options become available: – **Type** – Select **MD5** to use the MD5 hash algorithm for authentication (recommended), or select **Clear text** to use clear text for authentication. – **Key ID** – The identification number of the authentication key. This number must be shared with all other devices sending updates to and receiving updates from the security appliance. Valid values range from 1 to 255. – **Key** – The shared key used for authentication. This key must be shared with all other devices sending updates to and receiving updates from the security appliance. The key can be up to 16 characters.

RIP Page for PIX/ASA 7.2 and Later

Note From version 4.17, though Cisco Security Manager continues to support PIX and FWSM features/functionality, it does not support any bug fixes or enhancements.

Use this RIP page to enable and configure the Routing Information Protocol (RIP) on PIX and ASA devices running operating system 7.2 or later. The RIP page consists of these tabbed panels:

- RIP - Setup Tab, page 56-123
- RIP - Redistribution Tab, page 56-125
- RIP - Filtering Tab, page 56-126

- RIP - Interface Tab, page 56-127

Navigation Path

- (Device view) Select **Platform > Routing > RIP** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > RIP** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

  When creating a shared RIP policy, you must choose a Version in the Create a Policy dialog box, as follows:

  - **PIX/ASA 6.3-7.1 and FWSM**
  - **PIX/ASA 7.2 and Later**

  When assigning a shared RIP policy, be sure to assign the appropriate RIP policy for the device. For example, you cannot assign a PIX/ASA 7.2+ RIP policy to an FWSM.

Related Topics

- Configuring Static Routes, page 56-128
- Configuring OSPF, page 56-75
- Configuring No Proxy ARP, page 56-1
- RIP Page for PIX/ASA 6.3–7.1 and FWSM, page 56-120

RIP - Setup Tab

Use the Setup panel to define RIP on the security appliance, and to configure global RIP protocol parameters. You can only enable a single RIP process on the security appliance.

Navigation Path

You can access the Setup tab from the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Related Topics

- RIP - Redistribution Tab, page 56-125
- RIP - Filtering Tab, page 56-126
- RIP - Interface Tab, page 56-127
- Chapter 56, “Configuring Routing Policies on Firewall Devices”

Field Reference

Table 56-82 Setup Tab

Element	Description
Networks	Define one or more networks for RIP routing. Enter IP address(es), or enter or Select the desired Network/Hosts objects (see Understanding Networks/Hosts Objects, page 6-80); IP addresses must not contain any subnet information. There is no limit to the number of networks you can add to the security appliance configuration. The RIP routing updates will be sent and received only through interfaces on the specified networks. Also, if the network of an interface is not specified, the interface will not be advertised in any RIP updates.
Passive Interface	Use this option to specify passive interfaces on the security appliance, and by extension the active interfaces. The device listens for RIP routing broadcasts on passive interfaces, using that information to populate its routing tables, but does not broadcast routing updates on passive interfaces. Interfaces that are not designated as passive, receive and send updates. Choose one of these options: 1. **None** – No interfaces are designated as passive. 2. **All Interfaces** – All interfaces on the device are designated as passive, except those entered the Excluded Interfaces field below. 3. **Specified Interfaces** – Only those interfaces explicitly specified in the Interfaces field below are designated as passive.
Interfaces/Excluded Interfaces	Use this field to specify the interfaces excluded from the passive list, or those explicitly designated as passive, depending on your choice from the Passive Interface list above: • If you chose All Interfaces, this field is labeled **Excluded Interfaces**: enter or Select only those interfaces to be excluded (that is, those that are to be active not passive). • If you chose Specified Interfaces in the Passive Interface list, enter or Select those interfaces that are to be designated as passive. **Note** You cannot specify two different RIP configurations for the same interface.
RIP Version	Choose the RIP versions for sending and receiving RIP updates: • Receive Version 1 and 2, Send Version 1 • Send and Receive Version 1 • Send and Receive Version 2
Generate Default Route	When selected, a default route is generated for distribution, based on the Route Map you specify.
Route Map	Specify the route map to use for generating default routes. **Note** This field contains only the Route Map name. The Route Map is created and contained within a FlexConfig; see Chapter 7, "Managing FlexConfigs" for more information.

Table 56-82 Setup Tab (continued)

Element	Description
Enable Auto-Summary	When Send and Receive Version 2 is the chosen RIP Version, this option is available. When checked, automatic route summarization is enabled. Disable automatic summarization if you must perform routing between disconnected subnets. When automatic summarization is disabled, subnets are advertised. **Note** RIP Version 1 always uses automatic summarization—you cannot disable it.

RIP - Redistribution Tab

Use the Redistribution panel to manage redistribution routes. These are the routes that are being redistributed from other routing processes into the RIP routing process. See Add/Edit Redistribution Dialog Box, page 56-125 for more information.

Navigation Path

You can access the Redistribution tab from the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Related Topics

- RIP - Setup Tab, page 56-123
- RIP - Filtering Tab, page 56-126
- RIP - Interface Tab, page 56-127
- Chapter 56, “Configuring Routing Policies on Firewall Devices”

Add/Edit Redistribution Dialog Box

Use the Add Redistribution and Edit Redistribution dialog boxes to add and edit redistribution routes on the RIP - Redistribution Tab, page 56-125. These are the routes that are being redistributed from other routing processes into the RIP routing process. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Redistribution dialog boxes from the Redistribution tab on the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Field Reference

Table 56-83 Add/Edit Redistribution Dialog Box

Element	Description
Protocol to Redistribute	Choose the routing protocol to redistribute into the RIP routing process: • **Static** – Static routes. • **Connected** – Directly connected networks. • **OSPF** – Routes discovered by the OSPF routing process. If you choose OSPF, you must also enter the OSPF Process ID and, optionally, Match criteria.

Table 56-83 Add/Edit Redistribution Dialog Box (continued)

Element	Description
Process ID	Enter the process ID when the OSPF protocol is chosen.
Match	If you are redistributing OSPF routes into the RIP routing process, you can select specific types of OSPF routes to redistribute. Ctrl-click to select multiple types: • Internal – Routes internal to the autonomous system (AS) are redistributed. • External 1 – Type 1 routes external to the AS are redistributed. • External 2 – Type 2 routes external to the AS are redistributed. • NSSA External 1 – Type 1 routes external to a not-so-stubby area (NSSA) are redistributed. • NSSA External 2 – Type 2 routes external to an NSSA are redistributed. Match criteria are optional. The default is match Internal, External 1, and External 2.
Metric	The RIP metric type to apply to the redistributed routes. The two choices are: • **Transparent** – Use the current route metric. • **Specified Value** – Assign a specific metric value.
Metric Value	The metric value to be assigned; enter a value from 0 to 16.
Route Map	The name of a route map that must be satisfied before the route can be redistributed into the RIP routing process. **Note** This field contains only the route Map name. The contents of the route map are created and contained within a FlexConfig. See Chapter 7, "Managing FlexConfigs" for more information.

RIP - Filtering Tab

Use the Filtering panel to manage filters for the RIP policy. Filters are used to limit network information in incoming and outgoing RIP advertisements. See Add/Edit Filter Dialog Box, page 56-127 for more information.

Navigation Path

You can access the Filtering tab from the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Related Topics

- RIP - Setup Tab, page 56-123
- RIP - Redistribution Tab, page 56-125
- RIP - Interface Tab, page 56-127
- Chapter 56, "Configuring Routing Policies on Firewall Devices"

Add/Edit Filter Dialog Box

Use the Add Filter and Edit Filter dialog boxes to add and edit RIP filters on the RIP - Filtering Tab, page 56-126. Filters are used to limit network information in incoming and outgoing RIP advertisements. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Filter dialog boxes from the Filtering tab on the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Field Reference

Table 56-84 Add/Edit Filter Dialog Box

Element	Description
Traffic Direction	Choose the type of traffic to be filtered: **Inbound** or **Outbound**. **Note** If Traffic Direction is Inbound, you can define an Interface filter only.
Filter On	Specify whether the filter is based on an **Interface** or a **Route**. If you select Interface, enter or Select the name of the interface on which routing updates are to be filtered. If you select Route, choose the route type: • **Static** – Only static routes are filtered. • **Connected** – Only connected routes are filtered. • **OSPF** – Only OSPF routes discovered by the specified OSPF process are filtered. Enter the Process ID of the OSPF process to be filtered.
Filter ACLs	Enter or Select the name of one or more access control lists (ACLs) that define the networks to be allowed or removed from RIP route advertisements.

RIP - Interface Tab

Use the Interface panel to manage the interfaces configured to send and receive RIP broadcasts. See Add/Edit Interface Dialog Box, page 56-128 for more information.

Navigation Path

You can access the Interface tab from the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Related Topics

- RIP - Setup Tab, page 56-123
- RIP - Redistribution Tab, page 56-125
- RIP - Filtering Tab, page 56-126
- Chapter 56, “Configuring Routing Policies on Firewall Devices”

Add/Edit Interface Dialog Box

Use the Add Interface and Edit Interface dialog boxes to add and edit RIP interface configurations on the RIP - Interface Tab, page 56-127. Except for their titles, these two dialog boxes are identical.

Navigation Path

You can access the Add and Edit Interface dialog boxes from the Interface tab on the RIP Page for PIX/ASA 7.2 and Later, page 56-122.

Field Reference

Table 56-85 Add/Edit Interface Dialog Box

Element	Description
Interface	Enter or Select an interface defined on this appliance.
Send (Version)	These options let you override, for this interface, the global Send versions specified on the RIP - Setup Tab, page 56-123. Select the appropriate boxes to specify sending updates using RIP Version 1, Version 2, or both.
Receive (Version)	These options let you override the global Receive versions. Select the appropriate boxes to specify accepting updates using RIP Version 1 only, Version 2 only, or both.
Authentication type	Choose the authentication used on this interface for RIP broadcasts: • **None** – No authentication. • **MD5** – Employ MD5. • **Clear Text** – Employ clear-text authentication. If you choose MD5 or Clear Text, you must also provide the following authentication parameters: • **Key ID** – The ID of the authentication key. Valid values are from 0 to 255. • **Key** – The key used by the chosen authentication method. Can contain up to 16 characters. • **Confirm** – Enter the authentication key again, to confirm.

Configuring Static Routes

A static route is a specific path to a particular destination network that is manually defined on the current device. Static routes are used in a variety of situations, and can be a quick and effective way to route data from one network to another when there is no dynamic route to the destination, or when use of a dynamic routing protocol is not feasible.

All routes have a value or “metric” that represents its priority of use. (This metric is also referred to as “administrative distance.”) When two or more routes to the same destination are available, devices use administrative distance to decide which route to use.

For static routes, the default metric value is one, which gives them precedence over routes from dynamic routing protocols. If you increase the metric to a value greater than that of a dynamic route, the static route operates as a back-up in the event that dynamic routing fails. For example, Open Shortest Path First

(OSPF)-derived routes have a default administrative distance of 100. To configure a back-up static route that is overridden by an OSPF route, specify a metric value for the static route that is greater than 100. This is referred to as a "floating" static route.

There is a special kind of static route known as a default route, or a "zero-zero" route because all zeroes are used for both the destination address and subnet mask. The default static route serves as a catch-all gateway: if there are no matches for a particular destination in the device's routing table, the default route is used. The default route generally includes a next-hop IP address or local exit interface.

Use the Static Route page to maintain manually defined static routes. The Static Route table on this page lists all currently defined static routes, showing for each, the name of the interface or interface role for which the route is defined, the destination network(s), the next hop gateway, the route metric, whether the route is tunneled, and whether there is service-level agreement tracking for the route. For a detailed explanation of these fields, see Add/Edit Static Route Dialog Box, page 56-130 or Add/Edit IPv6 Static Route Dialog Box, page 56-131.

Static null0 Route Configuration

Typically ACLs are used for traffic filtering and they enable you to filter packets based on the information contained in their headers. In packet filtering, the ASA firewall examines packet headers to make a filtering decision, thus adding some overhead to the processing of the packets and affecting performance.

Static null 0 routing is a complementary solution to filtering. A static null0 route is used to forward unwanted or undesirable traffic into a black hole. The null interface null0, is used to create the black hole. Static routes are created for destinations that are not desirable, and the static route configuration points to the null interface. Any traffic that has a destination address that has a best match of the black hole static route is automatically dropped. Unlike with ACLs, static null0 routes do not cause any performance degradation.

The static null0 route configuration is used to prevent routing loops. BGP leverages the static null0 configuration for Remotely Triggered Black Hole routing.

Navigation Path

- (Device view) Select **Platform > Routing > Static Route** or **Platform > Routing > IPv6 Static Route** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Routing > Static Route** or **PIX/ASA/FWSM Platform > Routing > IPv6 Static Route** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 56, "Configuring Routing Policies on Firewall Devices"
- Add/Edit Static Route Dialog Box, page 56-130
- Add/Edit IPv6 Static Route Dialog Box, page 56-131
- Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Table Columns and Column Heading Features, page 1-49

Add/Edit Static Route Dialog Box

The Add/Edit Static Route dialog box lets you add or edit a static route.

Navigation Path

You can access the Add/Edit Static Route dialog box from the Static Routes page. Click the Add Row button to add a new static route; select an existing static route and click the Edit Row button to edit that route.

Related Topics

- Configuring Static Routes, page 56-128
- Chapter 56, "Configuring Routing Policies on Firewall Devices"

Field Reference

Table 56-86 Add/Edit Static Route Dialog Box

Element	Description
Interface	Enter or Select the interface to which this static route applies. Sending traffic to a Null0 interface results in dropping the packets destined to the specified network. This feature is useful in configuring Remotely Triggered Black Hole (RTBH) for BGP. For more information, see Configuring Static Routes, page 56-128. **Note** If Null0 is selected as the interface, the Gateway and Tunneled options are disabled.
Network	Enter or Select the destination network(s). You can provide one or more IP address/netmask entries, one or more Networks/Hosts objects, or a combination of both; separate the entries with commas. Enter "0.0.0.0/0" or "any" to specify a default route.
Gateway	Enter or Select the gateway router which is the next hop for this route. You can provide an IP address, or a Networks/Hosts object. **Note** If an IP address from one of the security appliance's interfaces is used as the Gateway IP address, the security appliance will resolve the designated IP address in the packet instead of resolving the Gateway IP address.
Metric	The Metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. Because the hop count includes the destination network, all directly connected networks have a metric of 1. Enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1. The maximum number of equal-cost (equal-metric) routes that can be defined per interface is three. You cannot add a route with the same metric on different interfaces that are on the same network.

Table 56-86 Add/Edit Static Route Dialog Box (continued)

Element	Description
Tunneled	Select this option to make this a tunnel route; can be used only for a default route. You can configure only one default tunneled gateway per device. The Tunneled option is not supported in transparent mode. Available only on PIX/ASA 7.0+ devices.
Route Tracking	To monitor route availability, enter or Select name of an SLA (service level agreement) object that defines the monitoring policy. Available only on PIX/ASA 7.2+ devices. For more information on route tracking, see Monitoring Service Level Agreements (SLAs) To Maintain Connectivity, page 51-8.

Add/Edit IPv6 Static Route Dialog Box

The Add/Edit IPv6 Static Route dialog box lets you add or edit an IPv6 static route. IPv6 static routes are only supported on the following devices:

- ASA 7.0 and later (Routed mode)
- ASA 8.2 and later (Transparent mode)
- FWSM 3.1 and later (Routed mode)

Navigation Path

You can access the Add/Edit IPv6 Static Route dialog box from the IPv6 Static Route page. Click the **Add Row** button to add a new static route; select an existing static route and click the **Edit Row** button to edit that route.

Related Topics

- Configuring Static Routes, page 56-128
- Chapter 56, "Configuring Routing Policies on Firewall Devices"

Field Reference

Table 56-87 Add/Edit IPv6 Static Route Dialog Box

Element	Description
Interface	Enter or Select the interface to which this static route applies.
IPv6 Network	Enter or Select the destination network(s). You can provide one or more IP address entries, one or more Networks/Hosts objects, or a combination of both; separate the entries with commas. Enter two colons (::) to specify a default route.
IPv6 Gateway	Enter or Select the gateway router which is the next hop for this route. You can provide an IP address, or a Networks/Hosts object. **Note** If an IP address from one of the security appliance's interfaces is used as the Gateway IP address, the security appliance will resolve the designated IP address in the packet instead of resolving the Gateway IP address.

Table 56-87 Add/Edit IPv6 Static Route Dialog Box (continued)

Element	Description
Metric	The Metric is a measurement of the "expense" of a route, based on the number of hops (hop count) to the network on which a specific host resides. Hop count is the number of networks that a network packet must traverse, including the destination network, before it reaches its final destination. Because the hop count includes the destination network, all directly connected networks have a metric of 1. Enter the number of hops to the destination network. Valid values range from 1 to 255; the default value is 1. The maximum number of equal-cost (equal-metric) routes that can be defined per interface is three. You cannot add a route with the same metric on different interfaces that are on the same network.
Tunneled	Select this option to specify the route as the default tunnel gateway for VPN traffic. You can configure only one default tunneled gateway per device. Available only on ASA 7.0+ devices in routed mode.

Configuring Policy Objects for ASA Routing Policies

There are several policy objects that you use with ASA routing policies. This reference explains the configuration of these policy objects.

This section contains the following topics:

- Understanding Route Map Objects, page 56-132
- Add or Edit Policy List Object Dialog Box, page 56-143
- Add or Edit Prefix List Object Dialog Box, page 56-146
- Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148
- Add or Edit As Path Object Dialog Boxes, page 56-151
- Add or Edit Community List Object Dialog Box, page 56-153
- Create BFD Template, page 56-70

Understanding Route Map Objects

You can use route maps to define the conditions for redistributing routes from one routing protocol into another, or to enable policy routing.

Route maps have many features in common with widely known ACLs. These are some of the traits common to both:

- They are an ordered sequence of individual statements, each has a permit or deny result. Evaluation of ACL or route maps consists of a list scan, in a predetermined order, and an evaluation of the criteria of each statement that matches. A list scan is aborted once the first statement match is found and an action associated with the statement match is performed.
- They are generic mechanisms—Criteria matches and match interpretation are dictated by the way that they are applied. The same route map applied to different tasks might be interpreted differently.

These are some of the differences between route maps and ACLs:

- Route maps frequently use ACLs as matching criteria.

  Note Route maps do not support ACLs that include a user, user group, security group tag, or fully qualified domain name objects.

- The main result from the evaluation of an ACL is a yes or no answer—An ACL either permits or denies input data. Applied to redistribution, an ACL determines if a particular route can (route matches ACLs permit statement) or can not (matches deny statement) be redistributed. Typical route maps not only permit (some) redistributed routes but also modify information associated with the route, when it is redistributed into another protocol.
- Route maps are more flexible than ACLs and can verify routes based on criteria which ACLs can not verify. For example, a route map can verify if the type of route is internal.
- Each ACL ends with an implicit deny statement, by design convention; there is no similar convention for route maps. If the end of a route map is reached during matching attempts, the result depends on the specific application of the route map. Fortunately, route maps that are applied to redistribution behave the same way as ACLs: if the route does not match any clause in a route map then the route redistribution is denied, as if the route map contained deny statement at the end.

Route maps are preferred if you intend to either modify route information during redistribution or if you need more powerful matching capability than an ACL can provide. If you simply need to selectively permit some routes based on their prefix or mask, we recommend that you use a route map to map to an ACL (or equivalent prefix list).

Note You must use a standard ACL as the match criterion for your route map. Using an extended ACL will not work, and your routes will never be redistributed. We recommend that you number clauses in intervals of 10 to reserve numbering space in case you need to insert clauses in the future.

Permit and Deny Clauses

Route maps can have permit and deny clauses. If the match criteria are met for this route map, and the permit keyword is specified, the route is redistributed as controlled by the set actions. If the match criteria are not met, and the permit keyword is specified, the next route map with the same map tag is tested. If a route passes none of the match criteria for the set of route maps sharing the same name, it is not redistributed by that set. If the match criteria are met for the route map and the deny keyword is specified, the route is not redistributed.

The following rules apply:

- If you use an ACL in a route map using a permit clause, routes that are permitted by the ACL are redistributed.
- If you use an ACL in a route map deny clause, routes that are permitted by the ACL are not redistributed.
- If you use an ACL in a route map permit or deny clause, and the ACL denies a route, then the route map clause match is not found and the next route-map clause is evaluated.

Match and Set Clause Values

Each entry in a route map statement contains a combination of match and set clauses. The match clause defines the criteria for whether appropriate packets meet the particular policy (that is, the conditions to be met). The set clause explains how the packets should be routed once they have met the match criteria.

For each route that is being redistributed, the router first evaluates the match criteria of a clause in the route map. If the match criteria succeed, then the route is redistributed or rejected as dictated by the permit or deny clause, and some of its attributes might be modified as defined by the set clause. If the match criteria fail, then this clause is not applicable to the route, and the software proceeds to evaluate the route against the next clause in the route map. Scanning of the route map continues until a clause is found whose match clause matches the route or until the end of the route map is reached.

A match or set value in each clause can be missed or repeated several times, if one of these conditions exists:

- If several Match Clause values are present in a clause, all must succeed for a given route in order for that route to match the clause (in other words, the logical AND algorithm is applied for multiple match commands).
- If a Match Clause value refers to several objects in one command, any of the objects should match (the logical OR algorithm is applied).
- If a Match Clause value is not present, all routes match the clause.
- If a Set Value is not present in a route map permit clause, then the route is redistributed without modification of its current attributes.

Note Do not configure a Set Value in a route map deny clause because the deny clause prohibits route redistribution—there is no information to modify.

A route map clause without a Match or Set value performs an action. An empty permit clause allows a redistribution of the remaining routes without modification. An empty deny clause does not allow a redistribution of other routes (this is the default action if a route map is completely scanned, but no explicit match is found).

BGP Match and BGP Set Clauses

In addition to the match and set values described above, BGP provides additional match and set capabilities to route maps.

The following route-map match clauses are supported with BGP:

- match AS path access list
- match community
- match policy list

The following route-map set clauses are supported with BGP:

- set AS path
- set community
- set automatic tag
- set local preference
- set weight
- set origin
- set next hop
- set IP prefix list

Creating and Using Route Map Objects

When configuring a policy that requires that you identify a route map, you can select or create route map objects by clicking the **Select** button next to the Route Map field. To create a new route map from the Route Map Object Selector dialog box, click the **Create** button beneath the route map list. You can also create route map objects from the Policy Object Manager by selecting **Route Map** from the Object Type Selector and then clicking the **New Object** button. For information on the specific fields available when creating a route map object, see Add or Edit Route Map Object Dialog Boxes, page 56-136.

Note about the use of Route Map Objects in BGP Policies

Some of the match and set criteria used in route maps are not supported in all BGP subcommands. For example:

The following route map match criteria:

- Match Clause tab > Match first hop interface of route, Match Next Hop (IPv4 and IPv6), Match Route Source (IPv4 and IPv6), Match Metric Route Value, and Match Tag
- BGP Match Clause tab > Match AS path access lists

and the following route map set criteria:

- Set Clause tab > Metric Values (all fields) and Metric Type
- BGP Set Clause tab > Set AS path, Prepend AS path, and Prepend last AS to the AS path

are not supported in the following places:

- BGP policy > IPv4 Address Family:
  - Aggregate Address tab > Attribute Map, Advertise Map, and Suppress Map
  - Neighbor tab > Filtering tab
  - Route Injection tab > Inject Map and Exist Map

Security Manager allows you to use route maps in your BGP configuration even if the route map contains unsupported match or set criteria and you will not receive a warning or error during validation. In such cases, deployment will fail and you will receive an error from the device in the following format:

```
…%"My-Route-map" used as BGP inbound route-map, nexthop match not supported….
```

Please refer to the ASA documentation for guidelines on the match/set criteria supported in route maps used in BGP configuration.

Related Topics

- Add or Edit Route Map Object Dialog Boxes, page 56-136
- Add or Edit Route Map Entry Dialog Box, page 56-137
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Add or Edit Route Map Object Dialog Boxes

Use the Add/Edit Route Map Object dialog box to create, copy and edit route map policy objects. You can use route maps to define the conditions for redistributing routes from one routing protocol into another, or to enable policy routing.

Navigation Path

Select **Manage > Policy Objects**, then select **Route Map** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Route Map Objects, page 56-132
- Add or Edit Route Map Entry Dialog Box, page 56-137
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-88 Add/Edit Route Map Object Dialog Box

Element	Description
Name	Enter a meaningful name for the route map object. The route map object name cannot be more than 58 characters. **Caution** Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.
Route Map table	The Route Map entries that are defined in the object. • To add a Route Map entry, click the **Add** button to open the Add or Edit Route Map Entry Dialog Box, page 56-137. • To edit a Route Map entry, select it and click the **Edit** button. • To delete a Route Map entry, select it and click the **Delete** button.

Table 56-88 Add/Edit Route Map Object Dialog Box (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Route Map Entry Dialog Box

Use the Add/Edit Route Map Entry dialog box to create a new route map entry for a Route Map object or to edit an existing one.

Navigation Path

From the Add or Edit Route Map Object Dialog Boxes, page 56-136, click the **Add** button beneath the Route Map table or select an entry in the table and click the **Edit** button.

Related Topics

- Understanding Route Map Objects, page 56-132
- Add or Edit Route Map Object Dialog Boxes, page 56-136
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-89 Add/Edit Route Map Entry Dialog Box

Element	Description
Sequence Number	A number, between 0 and 65535, that indicates the position a new route map entry will have in the list of route maps entries already configured for this route map object. **Tip** We recommend that you number clauses in intervals of at least 10 to reserve numbering space in case you need to insert clauses in the future.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Redistribution	Whether to redistribute a route or not. To allow redistribution for route matches, click **Permit**. To reject route matches from redistribution, select **Deny**. If you use an ACL in a route map Permit clause, routes that are permitted by the ACL are redistributed. If you use an ACL in a route map Deny clause, routes that are permitted by the ACL are not redistributed. In addition, if you use an ACL in a route map Permit or Deny clause, and the ACL denies a route, then the route map clause match is not found and the next route map clause is evaluated.
Match Clause Tab Select the Match Clause tab to choose routes to which this clause should be applied, and set the following parameters:	
Match first hop interface of route	Enable or disable matching routes that have their next hop out one of the interfaces specified. Enter or select the interfaces to match. Separate multiple entries with a comma. If you specify more than one interface, then the route can match either interface. Use the ellipsis to open the Interfaces Selector from which you can select one or more interfaces. You can also create new interface roles from the Interfaces Selector. For more information, see Understanding Interface Role Objects, page 6-73.
IPv4	
Match Address	Enable or disable matching of any routes that have a route address or match packet that is passed by one of the access lists specified. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.
Match Next Hop	Enable or disable matching of the next hop address of a route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Match Route Source	Enable or disable matching of the advertising source address of the route. For IPv4 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.
IPv6	
Match Address	Enable or disable matching of any routes that have a route address or match packet that is passed by one of the access lists specified. For IPv6 addresses, choose whether to use an access list or IPv6 Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148.
Match Next Hop	Enable or disable matching of the next hop address of a route. For IPv6 addresses, choose whether to use an access list or Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148.
Match Route Source	Enable or disable matching of the advertising source address of the route. For IPv6 addresses, choose whether to use an access list or IPv6 Prefix list for matching from the drop-down list and then enter or select the ACL objects or IPv6 Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or IPv6 Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Match Metric Route Value	Enable or disable matching the metric of a route. Type the metric values to use for matching in the Match Metric Route Value field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.
Match Tag	Enable or disable matching the security group tag of a route. Type the tag values to use for matching in the Match Tag field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.
Match Route Type	Enable or disable matching of the route type. Valid route types are External1, External2, Internal, Local, NSSA-External1, and NSSA-External2. When enabled, you can choose more than one route type from the list.
Set Clause Tab	
Select the Set Clause tab to modify the following information, which will be redistributed to the target protocol: **Note** You can specify just the Bandwidth value, all of the values, or none of the values.	
Bandwidth	Metric value or Bandwidth in Kbits per second; an integer value from 0 to 4294967295.
EIGRP Delay	EIGRP route delay, in tens of microseconds. Valid values range from 1 to 4294967295.
EIGRP Reliability	Likelihood of successful packet transmission for EIGRP expressed as a number from 0 to 255. The value 255 means 100 percent reliability; 0 means no reliability.
EIGRP Effective	Effective EIGRP bandwidth of a route expressed as a number from 1 to 255. The value 255 means 100 percent loading.
EIGRP MTU	Minimum MTU size of a route for EIGRP, in bytes. Valid values range from 1 to 4294967295.
Set Metric Type	Select to specify the type of metric for the destination routing protocol, and choose the metric type from the drop-down list: internal, type-1, or type-2.
BGP Match Clause Tab	
Match AS path access lists	Select to enable matching the BGP autonomous system path access list with the specified path access list. If you specify more than one path access list, then the route can match either path access list. Use the ellipsis to open the AS Path Object Selector from which you can select one or more AS path objects. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes, page 56-151.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Match community	Select to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Any route that does not match at least one Match community will not be advertised for outbound route maps. Use the ellipsis to open the Community List Object Selector from which you can select one or more Community List objects. You can also create new Community List objects from the Community List Object Selector. For more information, see Add or Edit Community List Object Dialog Box, page 56-153. To enable matching the BGP community exactly with the specified community, check the **Match the specified community exactly** check box.
Match policy list	Select to configure a route map to evaluate and process a BGP policy. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only. Use the ellipsis to open the Policy List Object Selector from which you can select one or more Policy List objects. You can also create new Policy List objects from the Policy List Object Selector. For more information, see Add or Edit Policy List Object Dialog Box, page 56-143.
BGP Set Clause Tab Select the BGP Set Clause tab to modify the following information, which will be redistributed to the BGP protocol:	
Set AS path	Select to modify an autonomous system path for BGP routes. • Select **Prepend AS path** to prepend an arbitrary autonomous system path string to BGP routes. Usually the local AS number is prepended multiple times, increasing the autonomous system path length. If you specify more than one AS path number then the route can prepend either AS number. • Select **Prepend last AS to the AS path** to prepend the AS path with the last AS number. Enter a value for the AS number from 1 to 10. • Select **Convert route tag into AS path** to convert the tag of a route into an autonomous system path.
Set community	Select to set the BGP communities attributes. • Select **None** to remove the community attribute from the prefixes that pass the route map. • Select **Specify community** to enter a community number, if applicable. Valid values are from 1 to 4294967295. Select **Add to the existing communities** to add the community to the already existing communities. • Select **Internet**, **no-advertise** or **no-export** to use one of the well-known communities.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Set Automatic-tag	Select to automatically compute the tag value.
Set local preference	Select to specify a preference value for the autonomous system path. Enter a value between 0 and 4294967295.
Set weight	Select to specify the BGP weight for the routing table. Enter a value between 0 and 65535.
Set origin	Select to specify the BGP origin code. Valid values are Local IGP and Incomplete.
Next hop IPv4	
Set next hop	Select to specify the output address of packets that fulfill the match clause of a route map: • Select **Specify IPv4 address** to enter the IPv4 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv4 address then the packets can output at either IP address. • Select **Use peer address** to set the next hop to be the BGP peer address.
Next hop IPv6	
Set next hop	Select to specify the output address of packets that fulfill the match clause of a route map: • Select **Specify IPv6 address** to enter the IPv6 address of the next hop to which packets are output. It need not be an adjacent router. If you specify more than one IPv6 address then the packets can output at either IP address. You can enter multiple values separated by commas. • Select **Use peer address** to set the next hop to be the BGP peer address.
Prefix List	
Set IPv4 prefix list	Select to set an IPv4 prefix list. Use the ellipsis to open the Prefix List Object Selector from which you can select one or more Prefix List objects. You can also create new Prefix List objects from the Prefix List Object Selector. For more information, see Add or Edit Prefix List Object Dialog Box, page 56-146.
Set IPv6 prefix list	Select to set an IPv6 prefix list. Use the ellipsis to open the Prefix List Object IPv6 Selector from which you can select one or more IPv6 Prefix List objects. You can also create new IPv6 Prefix List objects from the Prefix List Object Selector. For more information, see Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148.

Table 56-89 Add/Edit Route Map Entry Dialog Box (continued)

Element	Description
Policy Based Routing (PBR) Tab	
Click the Policy Based Routing tab to define policy for traffic flows, and lessening reliance on routes derived from routing protocols. PBR gives you more control over routing by extending and complementing the existing mechanisms provided by routing protocols. PBR allows you to set the IP precedence. It also allows you to specify a path for certain traffic, such as priority traffic over a high-cost link.	
Set Default Next-Hop IPv4 Address	Check the Set default next-hop IPv4 address check box to indicate where to output packets that pass a match clause of a route map for policy routing. In the IPv4 Address enter the destination address.
Set Default Next-Hop IPv6 Address	Check the Set default next-hop IPv6 address check box to indicate where to output packets that pass a match clause of a route map for policy routing. In the IPv6 Address enter the destination address.
Recursively find and set Next-Hop IPv4 Address	Check the Recursively find and set next-hop IP address check box and specify an IP address in the IPv4 Address field. In this case, the next-hop IP address need not be on a directly connected subnet.
Set Interfaces	Check the Set interfaces check box and select a destination interface from the Interfaces Selector dialog box.
Set Null0 Interfaces as Default Interface	Check the Set null0 interface as the default interface check box, if there is a need to completely black hole or drop some traffic.
Set do-not-fragment bit to either 0 or 1	Check the Set do-not-fragment bit to either 1or 0 and then select the appropriate radio button.
Set Differential Service Code point (DSCP) value in Q...	Check the Set differential service code point (DSCP) value in QoS bits for IPv4 packets check box and either enter a value between 0 and 63 or select a value from the Select Value drop-down list.
Set Differential Service Code point (DSCP) value in QOS bits for IPv6 packets	Check the Set differential service code point (DSCP) value in QoS bits for IPv6 packets check box and either enter a value between 0 and 63 or select a value from the Select Value drop-down list.

Add or Edit Policy List Object Dialog Box

Use the Add/Edit Policy List Object dialog box to create, copy, and edit policy list policy objects. You can create policy list objects to use when you are configuring route maps (see Understanding Route Map Objects, page 56-132).

When a policy list is referenced within a route map, all of the match statements within the policy list are evaluated and processed. Two or more policy lists can be configured with a route map. A policy list can also coexist with any other preexisting match and set statements that are configured within the same route map but outside of the policy list. When multiple policy lists perform matching within a route map entry, all policy lists match on the incoming attribute only.

Navigation Path

Select **Manage > Policy Objects**, then select **Policy List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Understanding Route Map Objects, page 56-132
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-90 Add/Edit Policy List Object Dialog Box

Element	Description
Name	The name of the object. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9. **Caution** Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.
Basic Tab	
Action	Whether to permit access for matching conditions or not. **Note** The Action for a policy list object cannot be changed after the initial creation of the policy list object.
Match Interface	Select to distribute routes that have their next hop out of one of the interfaces specified. Enter or select the interfaces to match. Separate multiple entries with a comma. If you specify more than one interface, then the route can match either interface. Use the ellipsis to open the Interfaces Selector from which you can select one or more interfaces. You can also create new interface roles from the Interfaces Selector. For more information, see Understanding Interface Role Objects, page 6-73.

Table 56-90 Add/Edit Policy List Object Dialog Box (continued)

Element	Description
Match Address	Select to redistribute any routes that have a destination address that is permitted by a standard access list or prefix list. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.
Match Next-Hop	Select to redistribute any routes that have a next hop router address passed by one of the access lists or prefix lists specified. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.
Match Route Source	Select to redistribute routes that have been advertised by routers and access servers at the address specified by the access lists or prefix list. Choose whether to use an Access List or Prefix List for matching from the drop-down list and then enter or select the ACL objects or Prefix list objects you want to use for matching. Use the ellipsis to open the Access Control List Object Selector or Prefix List Object Selector from which you can select one or more objects. You can also create new objects from the object selector. For more information, see Add or Edit Access List Dialog Boxes, page 6-59 or Add or Edit Prefix List Object Dialog Box, page 56-146.
Advanced Tab	
Match AS Path	Select to match a BGP autonomous system path. If you specify more than one AS path, then the route can match either AS path. Use the ellipsis to open the AS Path Object Selector from which you can select one or more AS path objects. You can also create new AS path objects from the AS Path Object Selector. For more information, see Add or Edit As Path Object Dialog Boxes, page 56-151.

Table 56-90 Add/Edit Policy List Object Dialog Box (continued)

Element	Description
Match Community Rules	Select to enable matching the BGP community with the specified community. If you specify more than one community, then the route can match either community. Use the ellipsis to open the Community List Object Selector from which you can select one or more Community List objects. You can also create new Community List objects from the Community List Object Selector. For more information, see Add or Edit Community List Object Dialog Box, page 56-153. To enable matching the BGP community exactly with the specified community, check the **exact-match** check box.
Match Metric	Enable or disable matching the metric of a route. Type the metric values to use for matching in the Match Metric field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified metric. The metric values can range from 0 to 4294967295.
Match Tag	Enable or disable matching the security group tag of a route. Type the tag values to use for matching in the Match Tag field. You can enter multiple values separated by commas. This setting allows you to match any routes that have a specified security group tag. The tag values can range from 0 to 4294967295.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Prefix List Object Dialog Box

Use the Add/Edit Prefix List Object dialog box to create, copy and edit prefix list policy objects. You can create prefix list objects to use when you are configuring route maps (see Understanding Route Map Objects, page 56-132), policy maps (see Add or Edit Policy List Object Dialog Box, page 56-143), OSPF Filtering (see Add/Edit Filtering Dialog Box, page 56-93), or BGP Neighbor Filtering (see Add/Edit Neighbor Dialog Box, page 56-11).

Area Border Router (ABR) type 3 link-state advertisement (LSA) filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number.

Navigation Path

Select **Manage > Policy Objects**, then select **Prefix List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Add or Edit Prefix List Entry Dialog Box, page 56-148
- Understanding Route Map Objects, page 56-132
- Add or Edit Policy List Object Dialog Box, page 56-143
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-91 Add/Edit Prefix List Object Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9. **Caution** Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.
Prefix List table	The prefix list entries that are defined in the object. • To add a prefix list entry, click the **Add** button to open the Add or Edit Prefix List Entry Dialog Box, page 56-148. • To edit a prefix list entry, select it and click the **Edit** button. • To delete a prefix list entry, select it and click the **Delete** button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Table 56-91 Add/Edit Prefix List Object Dialog Box (continued)

Element	Description
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Prefix List Entry Dialog Box

Use the Add/Edit Prefix List Entry dialog box to create a new prefix list entry or edit an existing one.

Navigation Path

From the Add or Edit Prefix List Object Dialog Box, page 56-146, click the **Add** button beneath the Prefix List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 56-92 Add/Edit Prefix List Entry Dialog Box

Element	Description
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Sequence No	(Optional) Unique number that indicates the position a new prefix list entry will have in the list of prefix list entries already configured for this object. If left blank, the sequence number will default to five more than the largest sequence number currently in use.
IP Address	Specify the prefix number in the format of IP address/mask length.
Minimum Prefix Length	(Optional) Enter the minimum prefix length. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.
Maximum Prefix Length	(Optional) Enter the maximum prefix length. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Add or Edit Prefix List IPv6 Object Dialog Box

Use the Add/Edit Prefix List IPv6 Object dialog box to create, copy and edit IPv6 prefix list policy objects. You can create IPv6 prefix list objects to use when you are configuring route maps (see Understanding Route Map Objects, page 56-132), policy maps (see Add or Edit Policy List Object Dialog Box, page 56-143), OSPF Filtering (see Add/Edit Filtering Dialog Box, page 56-93), or BGP Neighbor Filtering (see Add/Edit Neighbor Dialog Box, page 56-11).

Area Border Router (ABR) type 3 link-state advertisement (LSA) filtering extends the capability of an ABR that is running OSPF to filter type 3 LSAs between different OSPF areas. Once a prefix list is configured, only the specified prefixes are sent from one OSPF area to another OSPF area. All other prefixes are restricted to their OSPF area. You can apply this type of area filtering to traffic going into or coming out of an OSPF area, or to both the incoming and outgoing traffic for that area.

When multiple entries of a prefix list match a given prefix, the entry with the lowest sequence number is used. For efficiency, you may want to put the most common matches or denials near the top of the list by manually assigning them a lower sequence number.

Navigation Path

Select **Manage > Policy Objects**, then select **Prefix ListIPv6** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Add or Edit IPv6 Prefix List Entry Dialog Box, page 56-150
- Understanding Route Map Objects, page 56-132
- Add or Edit Policy List Object Dialog Box, page 56-143
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-93 Add/Edit IPv6 Prefix List Object Dialog Box

Element	Description
Name	The IPv6 Prefix List object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9. **Caution** Security Manager allows you to rename these objects even though you cannot rename them on the device. When you rename these objects in Security Manager, the name change is accomplished by negating the existing CLI, and then issuing new CLI to create and assign the object using the new name. This initial negation may cause routing/network issues in your environment. Security Manager will not provide a warning message about these consequences when you rename the object.
Description	An optional description of the object.

Table 56-93 Add/Edit IPv6 Prefix List Object Dialog Box (continued)

Element	Description
IPv6 Prefix List table	The IPv6 prefix list entries that are defined in the object. • To add an IPv6 prefix list entry, click the **Add** button to open the Add or Edit IPv6 Prefix List Entry Dialog Box, page 56-150. • To edit an IPv6 prefix list entry, select it and click the **Edit** button. • To delete an IPv6 prefix list entry, select it and click the **Delete** button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit IPv6 Prefix List Entry Dialog Box

Use the Add/Edit IPv6 Prefix List Entry dialog box to create a new IPv6 prefix list entry or edit an existing one.

Navigation Path

From the Add or Edit Prefix List IPv6 Object Dialog Box, page 56-148, click the **Add** button beneath the Prefix List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 56-94 Add/Edit Prefix List Entry Dialog Box

Element	Description
Action	Select the Permit or Deny radio button to indicate the redistribution access.
Sequence No	(Optional) Unique number that indicates the position a new IPv6 prefix list entry will have in the list of IPv6 prefix list entries already configured for this object. If left blank, the sequence number will default to five more than the largest sequence number currently in use. **Note** Sequence Number must be in the range of 1 to 4294967295
IPv6 Address	Specify the prefix number in the format: IPv6 address/mask length where mask length is less than or equal to 128.
Minimum Prefix Length	(Optional) Enter the minimum prefix length in the range of 1 to 128. The value must be greater than the mask length and less than or equal to the Maximum Prefix Length, if specified.

Table 56-94 Add/Edit Prefix List Entry Dialog Box (continued)

Element	Description
Maximum Prefix Length	(Optional) Enter the maximum prefix length in the range of 1 to 128. The value must be greater than or equal to the Minimum Prefix Length, if present, or greater than the mask length if the Minimum Prefix Length is not specified.

Add or Edit As Path Object Dialog Boxes

Use the Add/Edit As Path Object dialog box to create, copy and edit autonomous system (AS) path policy objects. You can create AS path objects to use when you are configuring route maps (see Understanding Route Map Objects, page 56-132), policy maps (see Add or Edit Policy List Object Dialog Box, page 56-143), or BGP Neighbor Filtering (see Add/Edit Neighbor Dialog Box, page 56-11).

An AS path filter allows you to filter the routing update message by using access lists and look at the individual prefixes within an update message. If a prefix within the update message matches the filter criteria then that individual prefix is filtered out or accepted depending on what action the filter entry has been configured to carry out.

Note AS path object names must be a unique integer from 1-500. If an AS path object is discovered from a device or configuration file that uses the same name as an existing AS path object, the AS path object on Security Manager will be overwritten regardless of the Allow Device Override for Discovered Policy Objects setting on the Security Manager Administration - Discovery page.

Navigation Path

Select **Manage > Policy Objects**, then select **As Path** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Add or Edit As Path Entry Dialog Box, page 56-152
- Understanding Route Map Objects, page 56-132
- Add or Edit Policy List Object Dialog Box, page 56-143
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-95 Add/Edit As Path Object Dialog Box

Element	Description
Name	Enter a name for the AS Path Filter. Specify a unique value between 1 and 500.
Description	An optional description of the object.
AS Path table	The AS path entries that are defined in the object. • To add an AS path entry, click the **Add** button to open the Add or Edit As Path Entry Dialog Box, page 56-152. • To edit an AS path entry, select it and click the **Edit** button. • To delete an AS path entry, select it and click the **Delete** button. • To rearrange the entries, select an entry and then click the **Move Up** or **Move Down** button.
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit As Path Entry Dialog Box

Use the Add/Edit As Path Entry dialog box to create a new autonomous system (AS) path entry or edit an existing one.

Navigation Path

From the Add or Edit As Path Object Dialog Boxes, page 56-151, click the **Add Row** button beneath the As Path table or select an entry and click the **Edit Row** button.

Field Reference

Table 56-96 Add/Edit As Path Entry Dialog Box

Element	Description
Action	Select the **Permit** or **Deny** radio button to indicate the redistribution access.
Reg Exp	Specify the regular expression that defines the AS path filter. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions, page 17-109.

Add or Edit Community List Object Dialog Box

Use the Add/Edit Community List Object dialog box to create, copy and edit community list policy objects. You can create community list objects to use when you are configuring route maps (see Understanding Route Map Objects, page 56-132) or policy maps (see Add or Edit Policy List Object Dialog Box, page 56-143).

A community is a group of destinations that share some common attribute. You can use community lists to create groups of communities to use in a match clause of a route map. Just like an access list, a series of community lists can be created. Statements are checked until a match is found. As soon as one statement is satisfied, the test is concluded.

Navigation Path

Select **Manage > Policy Objects**, then select **Community List** from the Object Type Selector. Right-click inside the work area and select **New Object** or right-click a row and select **Edit Object**.

Related Topics

- Add or Edit Community List Entry Dialog Box, page 56-154
- Understanding Route Map Objects, page 56-132
- Add or Edit Policy List Object Dialog Box, page 56-143
- Policy Object Manager, page 6-4
- Selecting Objects for Policies, page 6-2
- Creating Policy Objects, page 6-9
- Editing Objects, page 6-12
- Using Category Objects, page 6-13
- Managing Object Overrides, page 6-17
- Allowing a Policy Object to Be Overridden, page 6-18

Field Reference

Table 56-97 Add/Edit Community List Object Dialog Box

Element	Description
Name	The object name, which can be up to 128 characters. Object names are not case-sensitive. For more information, see Creating Policy Objects, page 6-9.
Description	An optional description of the object.
Community List table	The community list entries that are defined in the object. • To add a community list entry, click the **Add** button to open the Add or Edit Community List Entry Dialog Box, page 56-154. • To edit a community list entry, select it and click the **Edit** button. • To delete a community list entry, select it and click the **Delete** button. • To rearrange the entries, select an entry and then click the **Move Up** or **Move Down** button.

Table 56-97 Add/Edit Community List Object Dialog Box (continued)

Element	Description
Category	The category assigned to the object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.
Allow Value Override per Device Overrides Edit button	Whether to allow the object definition to be changed at the device level. For more information, see Allowing a Policy Object to Be Overridden, page 6-18 and Understanding Policy Object Overrides for Individual Devices, page 6-18. If you allow device overrides, you can click the **Edit** button to create, edit, and view the overrides. The **Overrides** field indicates the number of devices that have overrides for this object.

Add or Edit Community List Entry Dialog Box

Use the Add/Edit Community List Entry dialog box to create a new community list entry or edit an existing one.

Navigation Path

From the Add or Edit Community List Object Dialog Box, page 56-153, click the **Add** button beneath the Community List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 56-98 Add/Edit Community List Entry Dialog Box

Element	Description
Type	Select the **Standard** or **Expanded** radio button to indicate the community rule type. **Note** You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.
Action	Select the **Permit** or **Deny** radio button to indicate the redistribution access.
Communities	Specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
internet	Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
no-advertise	Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
no-export	Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.
Expressions	For an expanded community list, specify the regular expression. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions, page 17-109.

Add or Edit Community List Entry Dialog Box

Use the Add/Edit Community List Entry dialog box to create a new community list entry or edit an existing one.

Navigation Path

From the Add or Edit Community List Object Dialog Box, page 56-153, click the **Add** button beneath the Community List table or select an entry in the table and click the **Edit** button.

Field Reference

Table 56-99 Add/Edit Community List Entry Dialog Box

Element	Description
Type	Select the **Standard** or **Expanded** radio button to indicate the community rule type. **Note** You cannot have entries using Standard and entries using Expanded community rule types in the same Community List object.
Action	Select the **Permit** or **Deny** radio button to indicate the redistribution access.
Communities	Specify a community number. Valid values can be from 1 to 4294967295 or from 0:1 to 65534:65535.
internet	Select to specify the Internet well-known community. Routes with this community are advertised to all peers (internal and external).
no-advertise	Select to specify the no-advertise well-known community. Routes with this community are not advertised to any peer (internal or external).
no-export	Select to specify the no-export well-known community. Routes with this community are advertised to only peers in the same autonomous system or to only other sub-autonomous systems within a confederation. These routes are not advertised to external peers.
Expressions	For an expanded community list, specify the regular expression. For information on the metacharacters you can use to build regular expressions, see Metacharacters Used to Build Regular Expressions, page 17-109.

Configuring Security Policies on Firewall Devices

You can configure general security settings for the device using the General page and the Timeouts page under Platform > Security. You can enable anti-spoofing on interfaces, configure IP fragment settings, and configure a variety of timeout values for the device.

This chapter contains the following topics:

- General Page, page 57-1
- Configuring Timeouts, page 57-4

General Page

Use the General page to configure security settings that help protect against malformed packets, spoofed packets, fragmented packets, and denial of service attacks. See Configuring Floodguard, Anti-Spoofing and Fragment Settings, page 57-2 for more information about the settings on this page.

Navigation Path

- (Device view) Select **Platform > Security > General** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Security > General** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 57, "Configuring Security Policies on Firewall Devices"
- Add/Edit General Security Configuration Dialog Box, page 57-3
- Configuring Timeouts, page 57-4

Field Reference

Table 57-1 General Page

Element	Description
Disable Floodguard (PIX 6.3 and FWSM 2.x only)	Check this box to disable Floodguard on the firewall device. This option is available only on PIX 6.3 and FWSM 2.x devices. See Configuring Floodguard, Anti-Spoofing and Fragment Settings, page 57-2 for more information about the Floodguard feature.

Table 57-1 General Page (continued)

Element	Description
Global Fragment Settings	
Use these options to configure global fragment settings for the device. You can override these settings for individual interfaces; see Add/Edit General Security Configuration Dialog Box, page 57-3 for more information.	
Enable Default Settings	Check this box to enable the default fragment settings fields.
Size	Specify the maximum number of fragments that can be in the IP re-assembly database waiting for re-assembly. The default is 200.
Chain	Specify the maximum number of fragments into which a full IP packet can be fragmented. The default is 24 packets.
Timeout	Specify the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.
Interface Configuration Table	
This table lists all interfaces on which individual anti-spoofing and fragment settings have been defined. Refer to Configuring Floodguard, Anti-Spoofing and Fragment Settings, page 57-2 for more information about these settings. Refer to Add/Edit General Security Configuration Dialog Box, page 57-3 for more information about configuring these settings on individual interfaces.	

Configuring Floodguard, Anti-Spoofing and Fragment Settings

Use the General page under Platform > Security to enable or disable Floodguard (on a PIX 6.3 or FWSM 2.x device), to enable Unicast Reverse Path Forwarding (anti-spoofing) on individual interfaces, and to configure IP fragment settings for the device, and for each interface of the device.

Floodguard

Floodguard lets you reclaim firewall resources if the user authentication subsystem runs out of resources. If an inbound or outbound `uauth` connection is being attacked or overused, the firewall will actively reclaim TCP user resources.

If the user authentication subsystem is depleted, TCP user resources in different states are reclaimed in the following order, depending on urgency:

1. Timewait
2. LastAck
3. FinWait
4. Embryonic
5. Idle

Floodguard is enabled by default. This option applies only to PIX 6.3 or FWSM 2.x devices.

Anti-spoofing

Unicast Reverse Path Forwarding (RPF) guards against IP spoofing—a packet using an incorrect source IP address to obscure its true source—by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the security appliance looks only at the destination address when determining where to forward the packet. Unicast RPF instructs the security appliance to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the security appliance, the security appliance routing table must include a route back to the source address. See RFC 2267 for more information.

With outside traffic, for example, the security appliance can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the security appliance uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, the security appliance drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the security appliance drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route look-up. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fragment Settings

Fragment settings provide management of packet fragmentation and improve compatibility with the Network File System (NFS). By default, the security appliance allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to allow fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the security appliance, as fragmented packets are often used as DoS attacks.

Related Topics

- General Page, page 57-1
- Add/Edit General Security Configuration Dialog Box, page 57-3

Add/Edit General Security Configuration Dialog Box

Use the Add/Edit General Security Configuration dialog box to enable or disable anti-spoofing, and to configure override fragment settings, for an interface.

Navigation Path

You can access the Add/Edit General Security Configuration dialog box from the Anti-Spoofing and Fragment Interface Configuration table on the Platform > Security > General Page, page 57-1.

Related Topics

- Chapter 57, “Configuring Security Policies on Firewall Devices”

- Configuring Floodguard, Anti-Spoofing and Fragment Settings, page 57-2

Field Reference

Table 57-2 Add/Edit General Security Configuration Dialog Box

Element	Description
Interface	Enter or Select the name of the interface for which you want to configure anti-spoofing or fragment settings.
Enable Anti-Spoofing	Check this box to enable Unicast RPF (anti-spoofing) on the specified interface.
Override Default Fragment Settings	To override the default fragment settings on the specified interface, check this box to enable the following fields, and then enter the new values. See the General Page, page 57-1 for the default global fragment settings on the device.
Size	Specify the maximum number of fragments that can be in the IP re-assembly database waiting for re-assembly for the specified interface. The default is 200.
Chain	Specify the maximum number of fragments into which a full IP packet can be fragmented for the specified interface. The default is 24 packets.
Timeout	Specify the maximum number of seconds to wait for an entire fragmented packet to arrive on the specified interface. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

Configuring Timeouts

The Timeouts page lets you set a variety of timeout values on the security appliance. All times are in the format **hh:mm:ss**.

These values represent idle timeouts for the connection and translation slots for various protocols. If a slot has not been used for the idle time specified, the resource is returned to the free pool. TCP connection slots are freed approximately 60 seconds after a normal connection close sequence.

Warning **We recommend that you do not change these values unless advised to do so by Customer Support.**

Navigation Path

- (Device view) Select **Platform > Security > Timeouts** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Security > Timeouts** from the Policy Types selector. Select an existing policy from the Policies selector, or create a new one.

Related Topics

- Chapter 57, "Configuring Security Policies on Firewall Devices"

Field Reference

Table 57-3 Timeouts Page

Element	Description
To change the timeout value for a parameter, click the radio button to the left of the parameter entry to activate it, and then enter the new value in the parameter field. To reset any value to its default, click the related Default button. Clicking the Disable button, where provided, disables the timeout by setting its value to 0:00:00. Follow either of the procedures in the previous paragraph to re-enable a disabled value.	
Translation Slot (xlate)	Length of time idle until a translation slot is freed. This value must be at least 1 minute; the default is 3 hours. Enter 0:00:00 to disable this timeout.
Connection (conn)	Length of time idle until a connection slot is freed. This value must be at least 5 minutes; the default is 1 hour. Click Disable or enter 0:00:00 to disable this timeout.
Half-Closed	Length of time idle until a half-closed TCP connection is closed. For ASA 9.1.2 and later devices, the minimum is 30 seconds. For all other devices, the minimum is 5 minutes. The default is 10 minutes. Click Disable or enter 0:00:00 to disable this timeout.
UDP	Length of time idle until a UDP protocol connection is closed. This value must be at least 1 minute; the default is 2 minutes. Click Disable or enter 0:00:00 to disable this timeout.
SCTP	Length of time idle until a SCTP protocol connection is closed. This value must be at least 1 minute; the default is 2 minutes. Click Disable or enter 0:00:00 to disable this timeout.
Connection Holddown	Length of time idle until traffic is forwarded. This is the time for which the ASA waits before forwarding traffic, to avoid route flapping.This value must be at least 1 second; the default is 15 seconds. Click Disable or enter 0:00:00 to disable this timeout.
ICMP (PIX 7.x+, ASA, FWSM 3.x+)	Length of time idle after which general ICMP states are closed.
RPC/Sun RPC	Length of time idle until a SunRPC slot is freed. This value must be at least 1 minute; the default is 10 minutes. Click Disable or enter 0:00:00 to disable this timeout.
H.225	Length of time idle until an H.225 signaling connection is closed. The H.225 default timeout is 1 hour (01:00:00). Setting the value to 00:00:00 means never close the connection. To close a connection immediately after all calls are cleared, enter 1 second (0:00:01).
H.323	Length of time idle until an H.323 media connection is closed. The default is 5 minutes. Click Disable or enter 0:00:00 to disable this timeout.
MGCP	Length of time idle after which MGCP media ports are closed. The default is 5 minutes (0:05:00). Click Disable or enter 0:00:00 to disable this timeout.
MGCP PAT (PIX 7.x+, ASA, FWSM 3.x+)	Length of time idle after which an MGCP PAT translation is removed. The minimum time is 30 seconds; the default is 5 minutes (0:05:00). Click Disable or enter 0:00:00 to disable this timeout.

Table 57-3 Timeouts Page (continued)

Element	Description
SIP	Length of time idle until an SIP signaling port connection is closed. This value must be at least 5 minutes; the default is 30 minutes. Click Disable or enter 0:00:00 to disable this timeout.
SIP Media	Length of time idle until an SIP media port connection is closed. This value must be at least 1 minute; the default is 2 minutes. Click Disable or enter 0:00:00 to disable this timeout.
SIP Disconnect (PIX 6.3(5), PIX/ASA 7.2+, FWSM 3.2+)	Length of time idle after which a SIP session is deleted if the 200 OK is not received for a CANCEL or a BYE message. The minimum value is 0:00:01; the maximum value is 0:10:00. The default value is 0:02:00.
SIP Invite (PIX 6.3(5), PIX/ASA 7.2+, FWSM 3.2+)	Length of time idle after which pinholes for PROVISIONAL responses and media xlates will be closed. The minimum value is 0:01:00; the maximum value is 0:30:00. The default value is 0:03:00.
SIP Provisional Media (PIX/ASA 7.2(3)+)	The timeout value for SIP provisional media connections; must be a value between 0:01:00 and 1193:00:00. The default is 2 minutes.
Auth. (uath) Absolute	Length of time until the authentication cache times out and new connections must be re-authenticated. The system waits until a user starts a new connection to prompt for re-authentication. This time must be shorter than the Translation Slot value. Click Disable or enter 0:00:00 to disable caching and require re-re-authentication on every new connection. **Note** Do not set this value to 0:00:00 if passive FTP is used on the connections. **Note** If you set this value to 0:00:00; HTTPS authentication may not work. If a browser initiates multiple TCP connections to load a Web page after HTTPS authentication, the first connection is permitted through, but subsequent connections trigger authentication. As a result, users are continuously presented with an authentication page, even after successful authentication. To work around this, set the authentication absolute timeout to 1 second. However, this workaround opens a one-second window of opportunity that might allow non-authenticated users through the firewall if they are coming from the same source IP address.
Auth. (uath) Inactivity	Length of time idle until the authentication cache times out and users have to re-authenticate new connections. This duration must be shorter than the Translation Slot value.
IGP	The Cisco ASA supports Non-Stop Forwarding from software Version 9.3.1 and later for dynamic routing protocols— Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF). The length of convergence time of Open Shortest Path First (OSPF) is 70 seconds by default. Using this field, you can change the length of the convergence time. This value must be within the range of 10 seconds to 1 hour and 40 seconds. The default value for IGP is 0:01:10.

CHAPTER 58

Configuring Service Policy Rules on Firewall Devices

This section describes configuring service policy rules. Service policies provide a consistent and flexible way to configure certain security appliance features, including priority queuing, application inspection, and QoS (quality of service). For example, you can use a service policy to create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications.

Note Detailed information regarding the range and implementation of particular service policies can be found on cisco.com. For example, the following references may be helpful:

Using Modular Policy Framework

QoS Configuration and Monitoring

This chapter contains the following topics:

- About Service Policy Rules, page 58-1
- About TCP State Bypass, page 58-3
- Priority Queues Page, page 58-4
- Service Policy Rules Page, page 58-5
- Configuring Traffic Flow Objects, page 58-18
- Configuring TCP Maps, page 58-22

About Service Policy Rules

Service policy rules encompass these features:

- TCP and general connection settings (including TCP State Bypass; see About TCP State Bypass, page 58-3)
- Content security control (CSC)
- Application inspection
- Intrusion Prevention Services
- QoS queuing and policing

- ASA CX redirection (see About the ASA CX, page 58-17)
- ASA FirePOWER redirection
- User statistics for identity-based firewall policies

The configuration options for these features are presented on two pages in Security Manager—**Priority Queues** and **Rules**—accessed by navigating to Platform > Service Policy.

Priority Queuing

Priority queuing establishes two queues on an interface, a Low Latency Queuing (LLQ) priority queue and a "best effort" queue. This lets you prioritize latency-sensitive traffic like voice and video, so it is transmitted ahead of other traffic. Packets in the priority queue are always transmitted before packets in the best effort queue.

Because queues are not of infinite size, they can fill and overflow. When a queue is full, additional packets cannot get into the queue and are dropped. This is called "tail drop." To minimize tail drop, you can increase the queue buffer size. You can also fine-tune the maximum number of packets allowed into the transmit queue. These options let you control the latency and robustness of priority queuing.

Priority queuing is a Quality of Service (QoS) feature. In Security Manager, priority queue size and transmit queue size are managed on the Priority Queues Page, page 58-4, while establishment of priority queuing for a traffic class is an option on the QoS tab of the Service Policy (MPC) Rule Wizard, which is accessed from the Service Policy Rules Page, page 58-5.

Application Inspection and QoS

Some applications require special handling by the security appliance, and specific application inspection engines are provided for this purpose. Specifically, applications that embed IP addressing information in the user data packet, or open secondary channels on dynamically assigned ports require special inspection.

Application inspection is enabled by default for many protocols, while it is disabled for others. In many cases, you can change the port which the application inspection engine monitors for traffic.

Application inspection engines work with network address translation (NAT) to help identify the location of embedded addressing information. This allows NAT to translate these embedded addresses, and to update any checksum or other fields that are affected by the translation.

Service policy rules define how specific types of application inspection are applied to different types of traffic processed by the security appliance. You can apply rules to specific interfaces, or globally to every interface.

These rules provide a means to configure security appliance features in a manner similar to the Cisco IOS software quality-of-service (QoS) CLI. For example, with service policy rules you can include IP Precedence as one of the criteria to identify traffic for rate-limiting. You can also create a timeout configuration that is specific to a particular TCP application, as opposed to one that applies to all TCP applications.

Traffic match criteria are used to define the types of traffic to which you want to apply application inspection. For example, TCP traffic on port 23 might be classified as the Telnet traffic class. You then might use the traffic class to apply connection limits.

Multiple traffic match criteria can be assigned to a single interface, but a packet will only match the first criteria within a specific service policy rule.

About TCP State Bypass

By default, all traffic that enters an ASA or FWSM is inspected using the Adaptive Security Algorithm, and is either allowed through or dropped based on the security policy. The device maximizes its firewall performance by checking the state of each packet—to determine whether this a new connection, or an established connection—and assigning it to the session management path (if it is a new connection SYN packet), the fast path (if it is an established connection), or the control-plane path (for advanced inspection).

Note TCP State Bypass is available on FWSM 3.2+ and ASA 8.2+ devices only.

TCP packets that match existing connections in the fast path can pass through the appliance without every aspect of the security policy being rechecked. This feature maximizes performance. However, the method of establishing the session in the fast path using the SYN packet, and the checks that occur in the fast path (such as TCP sequence number), require that both outbound and inbound flows for a connection pass through the same device, which is not the case in asymmetric routing environments.

For example, assume a new connection is assigned to security device 1. The SYN packet goes through the session management path, and an entry for the connection is added to the fast path table. If subsequent packets of this connection go through device 1, then the packets match the entry in the fast path, and are passed through. But if subsequent packets go to device 2, where a SYN packet did not go through the session management path, there is no entry in the fast path for the connection, and the packets are dropped.

Thus, if you have asymmetric routing configured on upstream routers, and traffic alternates between two security devices, enable TCP state bypass for those specific traffic flows. TCP state bypass alters the way sessions are established in the fast path and disables the fast path checks. TCP traffic is then treated much as a UDP connection is treated: when a non-SYN packet matching the specified networks enters the security device, and there is not a fast path entry, then the packet goes through the session management path to establish a connection in the fast path. Once in the fast path, the traffic bypasses the fast path checks.

Unsupported Features

The following features are not supported when you enable TCP state bypass:

- Application inspection – Application inspection requires both inbound and outbound traffic to go through the same security device, so application inspection is not supported with TCP state bypass.
- AAA authenticated sessions – When a user authenticates with one security device, traffic returning via the other security device will be denied because the user did not authenticate with that device.
- TCP Intercept, Maximum Embryonic Connections limit, TCP sequence number randomization – If TCP state bypass is enabled, the device does not keep track of the state of the connection, so these features are not applicable.
- Cisco CSC SSM (Content Security and Control Security Services Module) – SSM and SSC functionality cannot be used with TCP state bypass.

Compatibility with NAT

Because the translation session is established separately for each security device, be sure to configure static NAT on both devices for TCP state bypass traffic; if you use dynamic NAT, the address chosen for the session on device 1 will differ from the address chosen for the session on device 2.

Related Topics

- About Service Policy Rules, page 58-1

Priority Queues Page

Priority queues let you define how traffic is prioritized in the network. You can define a series of filters based on packet characteristics to cause traffic to be placed in a higher or lower priority queue. The queue with the highest priority is serviced first until it is empty, then the lower queues are serviced in sequence.

In Security Manager, priority queue size and transmit queue size are managed on this page, while establishment of priority queuing for a traffic class is an option on the QoS tab of the Service Policy (MPC) Rule Wizard, which is accessed from the Service Policy Rules Page, page 58-5.

The Priority Queue Configuration dialog box is used to add and edit these queues. Refer to Priority Queue Configuration Dialog Box, page 58-4 for descriptions of the fields displayed in the Priority Queues table on this page.

Note Priority queuing is are not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

Navigation Path

- (Device view) Select **Platform > Service Policy > Priority Queues** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Service Policy > Priority Queues** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

Related Topics

- Chapter 58, “Configuring Service Policy Rules on Firewall Devices”
- Insert/Edit Service Policy (MPC) Rule Wizard, page 58-6
- About Service Policy Rules, page 58-1
- Understanding Queuing Parameters, page 66-4

Priority Queue Configuration Dialog Box

Use the Priority Queue Configuration dialog box to define and edit the priority queues on the Priority Queues page.

Note Priority queuing is are not available on Catalyst 6500 service modules (the Firewall Services Module and the Adaptive Security Appliance Service Module).

Navigation Path

You open the Priority Queue Configuration dialog box by clicking the Add Row or Edit Row buttons on the Priority Queues Page, page 58-4.

Related Topics

- Chapter 58, “Configuring Service Policy Rules on Firewall Devices”

- Insert/Edit Service Policy (MPC) Rule Wizard, page 58-6
- About Service Policy Rules, page 58-1
- Understanding Queuing Parameters, page 66-4

Field Reference

Table 58-1 Priority Queue Configuration Dialog Box

Element	Description
Interface Name	Specify the interface to which this rule applies; you can enter the interface name, or click Select to choose an available interface.
Queue Limit	Enter the maximum number of packets that can be queued in a priority queue before it drops data. This limit must be in the range of 0 through 2048 packets.
Transmission Ring Limit	Enter the maximum number of packets allowed into the transmit queue. This fine-tuning of the transmit queue can reduce latency and offer better performance through the transmit driver. On PIX devices, this value can range from 3 through 128 packets. On ASAs prior to version 7.2, this limit can be in the range 3 through 256 packets, while on ASAs running version 7.2 and higher, the value can be in the range 3 through 512 packets.

Service Policy Rules Page

Use the Service Policy Rules page to define new service policy rules, and to edit or delete existing service policy rules.

Configuring Service Policy Rules consists of three tasks:

1. **Configure a service policy**. Create a service policy and determine the interfaces to which the service policy applies. For more information, see Step 1. Configure a Service Policy, page 58-6.
2. **Configure the traffic class**. Specify the criteria you want to use to identify the traffic to which the service policy applies. For more information, see Step 2. Configure the traffic class, page 58-7.
3. **Configure the actions**. Specify the actions that should be taken to protect information or resources, or to perform QoS functions for the traffic specified in this service policy. For more information, see Step 3. Configure the MPC actions, page 58-8.

The three tasks are performed using the Insert/Edit Service Policy (MPC) Rule Wizard, page 58-6. Refer to the individual task topics for descriptions of the fields displayed in the Service Policy Rules table on this page.

Navigation Path

- (Device view) Select **Platform > Service Policy > Rules** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > Service Policy > Rules** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or create a new one.

ASA CX Auth Proxy Configuration

The CXSC Auth Proxy button below the Service Policy Rules table opens the Add/Edit CXSC Auth Proxy Configuration dialog box, which is described in ASA CX Auth Proxy Configuration, page 58-17.

The CXSC Auth Proxy button is available below the Service Policy Rules table only in Device view; it is not visible in Policy view.

Note Security Manager uses "CXSC" in some places to refer to an ASA CX Security Services Processor (SSP).

Related Topics

- Chapter 58, "Configuring Service Policy Rules on Firewall Devices"
- About Service Policy Rules, page 58-1
- Standard rules table topics:
  - Using Rules Tables, page 12-8
  - Filtering Tables, page 1-48
  - Table Columns and Column Heading Features, page 1-49

Insert/Edit Service Policy (MPC) Rule Wizard

Use the Insert/Edit Service Policy (MPC) Rule wizard to add and edit service policy rules on the Service Policy Rules page. The Insert/Edit Service Policy (MPC) Rule wizard presents the following steps:

- Step 1. Configure a Service Policy, page 58-6
- Step 2. Configure the traffic class, page 58-7
- Step 3. Configure the MPC actions, page 58-8

Note "MPC" refers to what is now called the Modular Policy Framework. Refer to Using Modular Policy Framework for additional information.

Navigation Path

Open the Insert/Edit Service Policy (MPC) Rule wizard by clicking the Add Row or Edit Row button on the Service Policy Rules Page, page 58-5.

Step 1. Configure a Service Policy

The first step in using the Insert/Edit Service Policy (MPC) Rule Wizard to configure a Service Policy Rule involves enabling the rule and specifying the interface(s) to which it is applied.

Navigation Path

Open the Insert/Edit Service Policy (MPC) Rule wizard by clicking the Add Row or Edit Row button on the Service Policy Rules Page, page 58-5.

Related Topics

- Step 2. Configure the traffic class, page 58-7
- Step 3. Configure the MPC actions, page 58-8

Table 58-2 Insert/Edit Service Policy (MPC) Rule Wizard—Step 1. Configure a Service Policy.

Element	Description
Enable The Current MPC Rule	Check this box to enable this service policy rule. You can deselect this option if you want to define the rule now, but not deploy it to the device until later.
Category	To assign the rule to a category, choose the category from the list. Categories can help you organize and identify rules and objects. For more information, see Using Category Objects, page 6-13.
Description	Optionally, enter a description for the service policy rule.
Global - Applies to All Interfaces	Select this option to apply the rule globally to all interfaces. This option is not compatible with matching traffic based on the source or destination IP address using an access list.
Interfaces	Select this option to apply the rule to a specific interface or group of interfaces (or interface roles), and then enter or Select the name of an interface or interface object. This selection is required if you want to match traffic based on the source or destination IP address using an access list. **Note** Interface-specific rules take precedence over the global service policy for a given feature. For example, if you have a global policy with FTP inspection, and an interface policy with TCP connection limits, then both FTP inspection and TCP connection limits are applied to the interface. However, if you have a global policy with FTP inspection, and an interface policy with FTP inspection, then only the interface policy FTP inspection is applied to that interface.

Step 2. Configure the traffic class

The second step in using the Insert/Edit Service Policy (MPC) Rule Wizard to configure a Service Policy Rule involves specifying the traffic class to which the rule is applied.

Specify the class to use to match traffic for this rule:

- **Use class-default As The Traffic Class**—Select this option to use the traffic class `class-default` for this service policy. The class-default traffic class matches all traffic.
- **Traffic Class**—Select this option to apply this rule to a specific traffic class. Enter the name of the previously defined traffic class, or click **Select** to select it from the Traffic Flows Selector.

  You also can define or edit a traffic flow "on the fly" by clicking the either Create or Edit buttons in the Traffic Flows Selector. (Traffic flows are also created and edited on the Traffic Flows page of the Policy Object Manager.) See Configuring Traffic Flow Objects, page 58-18 for more information.

Related Topics

- Step 1. Configure a Service Policy, page 58-6
- Step 3. Configure the MPC actions, page 58-8

Step 3. Configure the MPC actions

The third step in the Insert/Edit Service Policy (MPC) Rule Wizard involves specifying IPS, CXSC, FirePOWER, Connection Setting, QoS, CSC, User Statistics, ScanSafe Web Security, and NetFlow parameters for the rule; each set of parameters is presented on a separate tabbed panel.

Related Topics

- Step 1. Configure a Service Policy, page 58-6
- Step 2. Configure the traffic class, page 58-7

Field Reference

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
Intrusion Prevention tab	
Enable IPS for this Traffic	Enables or disables intrusion prevention for this traffic flow. When this box is checked, the other parameters on this panel are available. **Note** These parameters are applicable only on ASA 7.0+ devices that have an IPS module installed. See About IPS Modules on ASA Devices, page 58-15 for more information.
IPS Mode	Select the operating mode for intrusion prevention: • **Inline**—This mode places the IPS module directly in the traffic flow. No traffic that you identified for IPS inspection can continue through the ASA without first passing through, and being inspected by, the IPS module. This mode is the most secure because every packet identified for inspection is analyzed before being allowed through. Also, the IPS module can implement a blocking policy on a packet-by-packet basis. However, this mode can affect throughput. • **Promiscuous**—This mode sends a duplicate stream of traffic to the IPS module. This is less secure than Inline mode, but has little impact on traffic throughput. Unlike Inline mode, in Promiscuous mode the IPS module cannot drop the original packets, it can only block traffic by instructing the ASA to shun the traffic, or by resetting the connection on the appliance. Also, while the IPS module is analyzing the traffic, a small amount of traffic may pass through the ASA before the IPS module can shun it.
On IPS Card Failure	Specify the action to be taken if the IPS module becomes inoperable. Select either: • **Open**—Permits traffic if the module or card fails. • **Close**—Blocks traffic if the module or card fails.
Virtual Sensor	Text box in which you can view, edit, or remove the virtual sensor in the service policy that you are adding or editing

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
CXSC tab	
Note Security Manager uses "CXSC" in places to refer to an ASA CX Security Services Processor (SSP).	
Enable CXSC For This Traffic	Check this box to enable redirection of this traffic flow to an ASA CX installed in the ASA. When this box is checked, the other parameters on this panel are available. **Note** These parameters are applicable only on ASA 5585-X devices running version 8.4(4)+ and ASA 55xx-X devices running version 9.1(1)+ that have an ASA CX SSP installed.
On Context Security Card Failure	Specify the action to be taken if the ASA CX becomes inoperable. Select either: • **Open** – If the ASA CX fails for any reason, the ASA will continue to pass traffic that would otherwise be redirected to the ASA CX. • **Close** – If the ASA CX fails, the ASA will drop any traffic that would otherwise be redirected to the ASA CX.
Enable Auth Proxy	Check this box to enable the authentication proxy, which is required if you want to use active authentication in the identity policies on the ASA CX. If not checked, no authentication is performed. **Note** You can change the port used for authentication proxy; see ASA CX Auth Proxy Configuration, page 58-17 for more information.
FirePOWER tab	
Enable FirePOWER Card For This Traffic	Check this box to enable redirection of this traffic flow to an ASA FirePOWER module installed in the ASA. When this box is checked, the other parameters on this panel are available. **Note** These parameters are applicable only on ASA 55xx-X devices running version 9.2(1)+.
On FirePOWER Card Failure	Specify the action to be taken if the ASA FirePOWER module becomes inoperable. Select either: • **Open** – If the ASA FirePOWER module fails for any reason, the ASA will continue to pass traffic that would otherwise be redirected to the ASA FirePOWER module. • **Close** – If the ASA FirePOWER module fails, the ASA will drop any traffic that would otherwise be redirected to the ASA FirePOWER module.
Enable Monitor Only	Sets the module to monitor-only mode. In monitor-only mode, the module can process traffic for demonstration purposes, but then drops the traffic. You cannot use the traffic-forwarding interface or the device for production purposes.

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
Connection Settings tab	
Enable Connection Settings For This Traffic	Enables or disables connection settings for this traffic flow. When this box is checked, the other parameters on this panel become active. From the Connection Settings tab you can configure maximum connections, embryonic connections, timeouts, and TCP parameters.
Maximum Connections	You can specify the maximum number of TCP and UDP connections, and the maximum number of embryonic connections for this traffic flow: • **Maximum TCP & UDP Connections** – Specify the maximum number of simultaneous TCP and UDP connections for the entire subnet, up to 65,535, for ASA versions earlier than 8.4(5); for ASA 8.4(5) and later, the maximum number is 2,000,000. The default is zero for both protocols, which means the maximum possible connections are allowed. • **Maximum TCP & UDP Connections Per Client** – For ASA/PIX 7.1+ only; specify the maximum number of simultaneous TCP and UDP connections on a per client basis. For ASA 8.4(5) and later, the maximum number is 2,000,000. • **Maximum Embryonic Connections** – For ASA/PIX 7.0+ only; specify the maximum number of embryonic connections per host, up to 65,535, for ASA versions earlier than 8.4(5); for ASA 8.4(5) and later, the maximum number is 2,000,000. An embryonic connection is a connection request that has not finished the necessary handshake between source and destination. This limit enables the TCP Intercept feature. The default is zero, which means the maximum embryonic connections. TCP Intercept protects inside systems from a DoS attack perpetrated by flooding an interface with TCP SYN packets. When the embryonic limit has been surpassed, the TCP intercept feature intercepts TCP SYN packets from clients to servers on a higher security level. SYN cookies are used during the validation process and help minimize the amount of valid traffic being dropped. Thus, connection attempts from unreachable hosts will never reach the server. This feature is not applicable if TCP State Bypass is enabled. • **Maximum Embryonic Connections Per Client** – For ASA/PIX 7.1+ only; specify the maximum number of embryonic connections on a per client basis. For ASA 8.4(5) and later, the maximum number is 2,000,000. This feature is not applicable if TCP State Bypass is enabled.

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
Connection Timeouts	You can specify the following connection timeout settings for this traffic flow: • **Embryonic Connection Timeout** – Specify the idle time until an embryonic connection slot is freed. Enter 0:00:00 to disable timeout for the connection. The default is 20 seconds for FWSMs, and 30 seconds for ASA/PIX devices. • **Half Closed Connection Timeout** – Specify the idle time until a half-closed connection slot is freed. Enter 0:00:00 to disable timeout for the connection. For FWSMs, the default value is 20 seconds; the maximum value is 255 seconds (four minutes, 15 seconds). For ASA 9.1.2 and later devices, the minimum is 30 seconds. For all other ASA/PIX devices, the minimum is 5 minutes. The default is 10 minutes for all ASA/PIX devices. • **Idle Connection Timeout** – Specify the idle time until a connection slot is freed. Enter 0:00:00 to disable timeout for the connection. This duration must be at least 5 minutes. The default is 1 hour.
Reset Connection Upon Timeout	If selected, connections are reset after a timeout occurs. Available for ASA/PIX 7.0(4)+ only.
Detect Dead Connections	Enables the Dead Connection Detection feature; available for ASA/PIX 7.2+ devices. Selecting this option enables these two fields: • **Dead Connection Detection Timeout** – Specify the period of time between retries when a dead connection is detected. The default is 15 seconds. • **Dead Connection Detection Retries** – Specify the number of retries to be performed after a dead connection is detected. The default is five.
Traffic Flow Idle Timeout	Specify the period of time between a traffic flow becoming idle and the flow's disconnection. Applicable to FWSM 3.2+ only. The default is 1 hour.
Enable TCP Normalization	Enables TCP normalization, and activates the TCP Map selection option. Applies to ASA/PIX 7.0+ only; not applicable if TCP State Bypass is enabled.
TCP map	Specify the TCP map to use for TCP normalization: enter or Select the name of a TCP map. For more information, see Configuring TCP Maps, page 58-22.

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
Randomize TCP Sequence Number	Enables the Randomize Sequence Number feature. Disable this feature only if another inline security appliance is also randomizing sequence numbers and the result is scrambling the data. Each TCP connection has two Initial Sequence Numbers: one generated by the client and one generated by the server. The security appliance randomizes the ISN that is generated by the host/server on the higher security interface. At least one ISN must be randomly generated so that attackers cannot predict the next ISN and potentially hijack the session. Not applicable if TCP State Bypass is enabled.
Enable TCP State Bypass	Enables TCP state bypass for this traffic flow. This is allows specific traffic flows in asymmetric routing environments when both the outbound and inbound flow for a connection do not pass through the same device. Applicable to FWSM 3.2+ and ASA 8.2+ only. See About TCP State Bypass, page 58-3 for more information.
Enable SCTP State Bypass (ASA 9.5.2 + only)	You can bypass Stream Control Transmission Protocol (SCTP) stateful inspection if you do not want SCTP protocol validation.
Enable Decrement TTL	Select this option to turn on decrementing of the time-to-live (TTL) value in packets passed by the security appliance. Applicable to PIX/ASA 7.2.2+ only.
Configure Flow Offload (For Firepower 9000/4000 series ASA 9.6(1) and above)	**Note** You must enable flow offload manually on the ASA and restart the device, before configuring flow offload in the Service Policy Wizard in Cisco Security Manager. Flow offload and flow offload statistics are supported on the ASA only in the single context and system context modes. It is not supported in the admin or user context. ASA supports flow offload starting from version 9.5.2(1); however Cisco Security Manager supports flow offload from ASA 9.6(1). Select this option to offload specific traffic to a super-fast path; traffic is switched and processed in the NIC instead of the ASA. Offloading can help you improve performance for data-intensive applications such as large file transfers. **Tip** You can configure flow offload only when TCP State Bypass and SCTP State bypass are not enabled on your device.
QoS tab	
Enable QoS For This Traffic	Enables Quality of Service (QoS) options for this traffic flow. When selected, the Enable Priority For This Flow and the Traffic Policing options become active. **Note** The options on this tab are applicable to PIX/ASA 7.0+ devices only.
Enable Priority For This Flow	Enables strict scheduling priority for this flow. The priority queues must be defined on the Priority Queues Page, page 58-4.

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
Traffic Policing	Enables output and input traffic policing. Traffic policing lets you control the maximum rate of traffic transmitted or received on an interface.
Output (Traffic Policing)	Enables policing of traffic flowing out of the device. If you enable policing, you can specify the following values: • **Committed Rate** – The rate limit for this traffic flow; this is a value in the range 8,000 to 2,000,000,000, specifying the maximum speed (bits per second) allowed. • **Burst Rate** – A value in the range 1,000 to 512,000,000 that specifies the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value. • **Conform Action** – The action to take when the rate is less than the conform-burst value. Choices are Transmit or Drop. • **Exceed Action** – Take this action when the rate is between the conform-rate value and the conform-burst value. Choices are Transmit or Drop.
Input (Traffic Policing)	Enables policing of traffic flowing into the device; these options apply to ASA/PIX 7.2+ devices only. If you enable policing, you can specify the following values: • **Committed Rate** – The rate limit for this traffic flow; this is a value in the range 8,000 to 2,000,000,000, specifying the maximum speed (bits per second) allowed. • **Burst Rate** – A value in the range 1,000 to 512,000,000 that specifies the maximum number of instantaneous bytes allowed in a sustained burst before throttling to the conforming rate value. • **Conform Action** – The action to take when the rate is less than the conform-burst value. Choices are Transmit or Drop. • **Exceed Action** – Take this action when the rate is between the conform-rate value and the conform-burst value. Choices are Transmit or Drop.
CSC tab	
Enable Content Security Control For This Traffic	Enables or disables the use of the Cisco CSC SSM (Content Security and Control Security Services Module) for this traffic flow. When this box is checked, the On CSC SSM Failure options become available. These options are applicable on ASA 7.1+ devices only; they are not applicable if TCP State Bypass is enabled. The CSC SSM provides protection against viruses, spyware, spam, and other unwanted traffic by scanning the FTP, HTTP, POP3, and SMTP packets.
On CSC SSM Failure	Specifies the action to take if the CSC SSM becomes inoperable: • **Open** – Permits traffic if the CSC SSM fails. • **Close** – Blocks traffic if the CSC SSM fails.

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
User Statistics tab	
Enable user statistics accounting (ASA 8.4(2)+ only)	Whether to collect user statistics accounting information for identity-based firewall policies. These statistics are kept for users to which a firewall policy is applied based on user name or user group membership. Select the type of information you want to collect: • **Account for sent drop count** • **Account for sent packet, sent drop and received packet count**
Protocol Inspection tab	
Enable Scansafe Web Security for this traffic (ASA 9.0+ only)	Enables or disables the use of ScanSafe Web Security for this traffic flow. When this box is checked, two options become available: These options are applicable on ASA 9.0+ devices only. • **ScanSafe Policy Map**– enables policy map selection. • **On ScanSafe Tower Communication Failure**– specifies action the system should take if ScanSafe Tower communication fails.
Enable SCTP for this traffic (ASA 9.5.2 + only	Enables or disables the use of SCTP for this traffic flow. • **SCTP Policy Map**– enables policy map selection
Enable Diameter Inspection for this traffic (ASA 9.5.2 + only)	Enables or disables the use of Diameter inspection for this traffic flow. • **Diameter Policy Map**– enables policy map selection When Diameter Inspection is enabled, you can further enable inspection of encrypted traffic by selecting the **Enable encrypted traffic inspection** check box. You must select the **TLS Proxy** to be used for this inspection.
Enable LISP for this traffic (ASA 9.5.2 + only)	Enables or disables the use of LISP Inspection for this traffic flow. • **LISP Policy Map**– enables policy map selection
Enable Flow LISP mobility for devices (ASA 9.5.2 + only)	Enables flow mobility in clustering.
Enable STUN Inspection support for devices (ASA 9.6.2 + only)	Enables or disables the use of STUN inspection for this traffic flow. It is supported on ASA 9.6.2 and above in the single and multi-context mode. **Note** When you enable STUN inspection on the default inspection class, TCP/UDP port 3478 is watched for STUN traffic. The inspection supports IPv4 addresses and TCP/UDP only.STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, as the transaction ID is not replicated among units, when a unit fails after receiving a STUN Request and another unit received the STUN Response, the STUN Response will be dropped.
Enable M3UA for this traffic (ASA 9.6.2 + only)	Enables or disables the use of M3UA for this traffic flow. • **M3UA Policy Map**– enables policy map selection

Table 58-3 Insert/Edit Service Policy (MPC) Rule Wizard—Step 3. Configure the actions.

Element	Description
NetFlow tab	
Enable NetFlow for this traffic	Enables or disables the use of NetFlow for this traffic flow. When this box is checked, the NetFlow options become available.
Collectors	Specify the collectors that should be used when sending NetFlow events of a specific event type: **Note** Only use collectors that have been configured on the NetFlow page at Platform > Logging > NetFlow. • Flow Create Event • Flow Deny Event • Flow Tear Event • All Event Types **Note** Cisco Security Manager does not allow duplicate netflow collectors for ASA 9.6(4) to 9.7.0, and 9.8(2) and above devices. Ensure that you remove the duplicate collectors.

About IPS Modules on ASA Devices

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

You can install a variety of IPS modules, such as the Advanced Inspection and Prevention Security Services Module (AIP-SSM), in some ASA device models. The IPS modules supported by each ASA model differ. The IPS modules run advanced IPS software that provides proactive, full-featured intrusion prevention services to stop malicious traffic, including worms and network viruses, before they can affect your network.

The ASA IPS module runs separately from the adaptive security appliance, and you need to add it to the device inventory as a separate device. It is, however, integrated into the ASA traffic flow.

When you configure the ASA IPS module, you need to configure the service policy rules on the host ASA as well as the IPS policies on the IPS module. The service policy rules determine which traffic is inspected by the IPS module. For an overview of IPS policy configuration, see Overview of IPS Configuration, page 36-5.

When you identify traffic for IPS inspection, the traffic flows through the ASA and the IPS module as follows:

1. Traffic enters the ASA.
2. Firewall policies, such as interface access rules, are applied.
3. Traffic is sent to the IPS module over the backplane when you operate in inline mode. If you configure the system to use promiscuous mode, a copy of the traffic is sent to the IPS module.

   See **IPS Mode** in the Intrusion Prevention section of the Insert/Edit Service Policy (MPC) Rule wizard (Step 3. Configure the MPC actions, page 58-8) for more information about Inline and Promiscuous modes.

4. The IPS module applies its security policy to the traffic and takes appropriate actions.
5. Allowed traffic is sent back to the adaptive security appliance over the backplane. In Inline mode, the IPS module may block some traffic according to its security policy; in other words, that traffic is not passed back.
6. VPN policies are applied (if configured).
7. Traffic exits the ASA.

The following illustration depicts traffic flow when running the IPS module in Inline mode. In this example, the IPS module automatically blocks traffic that it identifies as an attack. All other traffic is returned to the ASA.

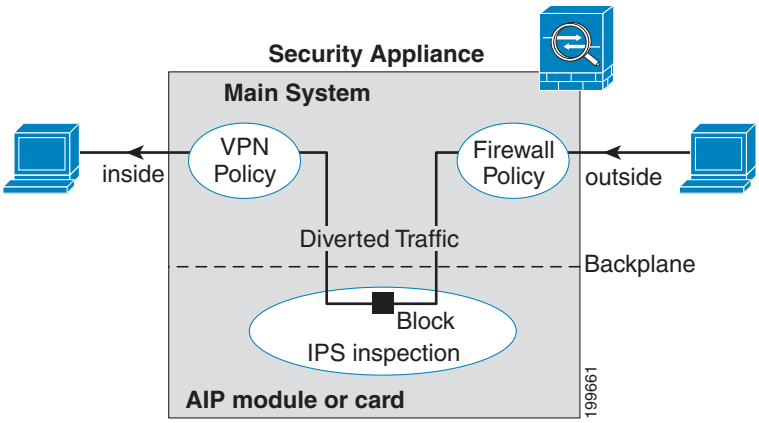

The next illustration depicts traffic flow when the IPS module is running in Promiscuous mode. In this example, the IPS module sends a shun message to the ASA for traffic it has identified as a threat.

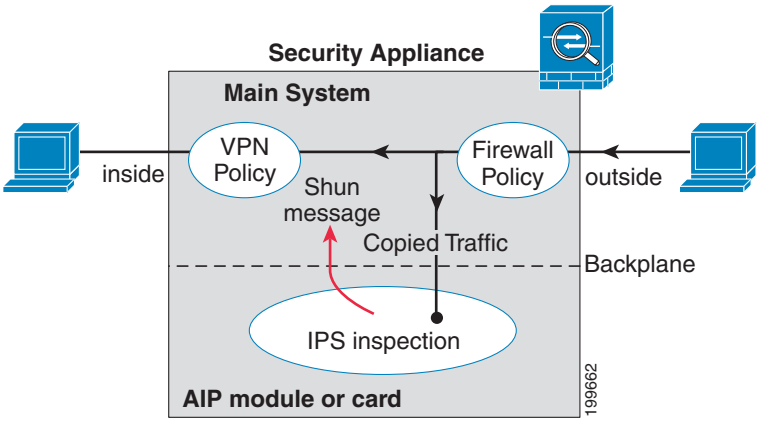

Related Topics

- Adding Devices to the Device Inventory, page 3-7

About the ASA CX

The ASA CX is a Security Services Processor (SSP) that can be installed in the Cisco ASA-5585-X series Adaptive Security Appliance. You configure the parent ASA to redirect traffic to the ASA CX, which then applies its security policies and either drops the traffic, or returns it to the ASA for further processing and routing to the next destination.

There are two basic policies that may need adjustment in the ASA when you add a ASA CX: access rules and inspection rules:

- Access rules, whether global rules or those applied to specific interfaces, are always applied before traffic is redirected to the ASA CX. Thus, the security card sees only the traffic already permitted, and does not process traffic that was dropped at entry to the ASA. Consider adjusting the rules to ensure that all traffic that you want the ASA CX to process is permitted.
- Inspection rules determine which traffic is inspected. The ASA CX does not inspect traffic that has already been inspected by the ASA. Therefore, you must ensure that you do not inspect traffic that you intend the ASA CX to inspect. Specifically, do not inspect HTTP traffic, because HTTP inspection is one of the core functions of the ASA CX. The default inspection rules on the ASA do not include HTTP inspection, so you must alter your inspection rules only if you added HTTP rules.

Determine if you need to create access rules for an interface, or global access rules that apply to all interfaces. Use the ASA access rules to filter traffic before it is redirected to the ASA CX. If you know there are classes of traffic that you never want to pass, it is more efficient to drop them immediately upon entry to the ASA.

If you have already established access rules, there is no requirement to change them. However, you should evaluate whether they need to be relaxed in order to have the ASA CX process certain types of traffic that you are now dropping by means of access rules.

Enabling traffic redirection to an installed ASA CX is described in Step 3. Configure the MPC actions, page 58-8 of the Insert/Edit Service Policy (MPC) Rule Wizard, page 58-6.

Related Topics

- About Service Policy Rules, page 58-1

ASA CX Auth Proxy Configuration

If you enabled ASA CX authentication proxy—on the **CXSC** tab during Step 3 of the Insert/Edit Service Policy (MPC) Rule Wizard; see Step 3. Configure the MPC actions, page 58-8—and you want to use a non-default port for active authentication, use the Add/Edit CXSC Auth Proxy Configuration dialog box to change the ASA CX Auth Proxy Port number.

If users must be prompted for authentication credentials, the prompting is done through this port.

Note Security Manager uses “CXSC” in some places to refer to an ASA CX Security Services Processor (SSP).

Navigation Path

Open the Add/Edit CXSC Auth Proxy Configuration dialog box by clicking the **CXSC Auth Proxy** button below the rules table on the Service Policy Rules Page, page 58-5.

Note The CXSC Auth Proxy button is available below the IPS, QoS, and Connection Rules table only in Device view; it is not visible in Policy view.

Related Topics

- Service Policy Rules Page, page 58-5

Field Reference

Table 58-4 Add/Edit CXSC Auth Proxy Configuration Dialog Box

Element	Description
CXSC Auth Proxy Port	The default authentication proxy TCP port is 885; however, if you change it, you must enter a port number between 1024 and 65535.

Configuring Traffic Flow Objects

Use the Add and Edit Traffic Flow dialog boxes to configure traffic-match definitions. These traffic-flow definitions correspond to class maps (the **class map** command) in the IPS, QoS and Connection Rules service policy for devices running the PIX 7.0+, ASA 7.0+, and FWSM 3.2+ operating systems. For more information on configuring these rules, see Chapter 58, "Configuring Service Policy Rules on Firewall Devices".

Navigation Path

Select **Manage > Policy Objects**, then select **Traffic Flows** from the Object Type selector. Right-click inside the work area and choose New Object, or right-click a row and choose Edit Object.

These dialog boxes also can be opened by clicking the Create or Edit buttons in the Traffic Flows Selector while defining a Service Policy rule. See for Step 2. Configure the traffic class, page 58-7 more information about selecting a Traffic Flow class.

Related Topics

- Creating Access Control List Objects, page 6-53

Field Reference

Table 58-5 Add and Edit Traffic Flow Dialog Boxes

Element	Description
Name	The name of the Traffic Flow object. A maximum of 40 characters is allowed. The name space for class maps is local to a security context. Therefore, the same name may be used in multiple security contexts. The maximum number of class maps per security context is 255.
Description	A description of the Traffic Flow (optional). A maximum of 1024 characters is allowed.

Table 58-5 Add and Edit Traffic Flow Dialog Boxes (continued)

Element	Description
Traffic Match Type	The type of traffic to match. The option you choose may change the fields in the dialog box; all possible fields are explained later in this table. The Traffic Match Type options are: • Any Traffic – Matches all traffic. • Source and Destination IP Address (access-list) – Matches the source and destination IP addresses in a packet based on the access control list that you specify. For ASA 8.4(2+) devices, the ACL can include FQDN objects and user specifications to perform identity-based traffic matching. • Default Inspection Traffic – Matches default inspection traffic. For a list of default settings, see Default Inspection Traffic, page 58-20. • Default Inspection Traffic with access list – Matches default inspection traffic limited by the access control list that you specify. • TCP or UDP or SCTP Destination Port – Matches traffic to the TCP or UDP or SCTP destination port or port range that you specify. Valid port numbers here are 0 to 65535. • RTP Range – Matches traffic to the range of UDP destination ports that you specify. Valid port numbers here are 2000 to 65535. • Tunnel Group – Matches the destination address based on flows in VPN tunnels belonging to a specified tunnel group. • IP Precedence Bits – Matches precedence values assigned to the traffic packets. You can select a maximum of four values. • IP DiffServe Code Points (DSCP) Values – Matches DSCP values associated with the traffic packets. You can select a maximum of eight values.
Variable Fields The following fields may appear in the Add and Edit Traffic Flow dialog boxes, depending on your choice in the Traffic Match Type field. This list is the complete set of possible fields.	
Available ACLs	A list of the access control list (ACL) objects that you can select for the map. Select the ACL that defines the target traffic, or click the Create button to add a new object. You can also select an object and click Edit to change its definition. If the list of objects is large, use the Filter field to limit the display (see Filtering Items in Selectors, page 1-45).
TCP or UDP or SCTP TCP/UDP/SCTP Port or Port Range	Radio buttons used to specify a protocol (either TCP, UDP, or SCTP), and a text field used to specify a destination port number or range of numbers to use when matching traffic based on the specified protocol/ports. You can specify a single port value, or a range of port numbers (for example, 0-2000). Valid port numbers are 0 to 65535.

Table 58-5 Add and Edit Traffic Flow Dialog Boxes (continued)

Element	Description
RTP Port Range	The range of RTP destination ports associated with the traffic flow. You must enter a port range within the valid range of 2000 to 65535. **Note** When you close the dialog box, the port range you entered is converted to port-span values by subtracting the start value from the end value. For example, if you enter the range 2001-3000 in the dialog box, "RTP port 2001 range 999" appears in the Match Value column of the Traffic Flows policy object table. Port-span values are expected by the device.
Tunnel group name Match Flow IP Destination Address	Lists available VPN tunnel groups. Choose one or enter the name of a group. You can also select Match Flow IP Destination Address to recognize the destination address as the match type. **Tip** You can use FlexConfig objects and policies to define a VPN tunnel group on a PIX 7.0+ device. For more information, see Understanding FlexConfig Policies and Policy Objects, page 7-2.
Available IP Precedence Match on IP Precedence	The IP precedence numbers. Select the values you want to match and click >> to add them to the Match table. Ctrl-click to select multiple values. You can select a maximum of four values. To remove a value from the Match table, select it and click <<.
Available DSCP Values Match on DSCP	The IP DiffServe Code Points (DSCP) numbers. Select the values you want to match and click >> to add them to the Match table. Ctrl-click to select multiple values. You can select a maximum of eight values. To remove a value from the Match table, select it and click <<.
Category	The category assigned to the traffic-flow object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Default Inspection Traffic

When you create a Traffic Flow policy object, you can choose to match the default inspection traffic. For more information, see Configuring Traffic Flow Objects, page 58-18. The following table lists the types of traffic included in the Default Inspection Traffic category.

Table 58-6 Default Inspection Traffic

Value	Port	NAT Limitations	Comments
CTIQBE	TCP/2748		
CuSeeMe	UDP/7648		
DNS over UDP	UDP/53	No NAT support for name resolution through WINS.	No PTR records are changed.
FTP	TCP/21		
GTP	UDP/2123, 3386		

Table 58-6 Default Inspection Traffic (continued)

Value	Port	NAT Limitations	Comments
H.323, H.225	TCP/1720, 1718	No NAT on same security interfaces. No static PAT.	
RAS	UDP/1718, 1719	No NAT on same security interfaces. No static PAT.	
HTTP	TCP/80		
ICMP	—		All ICMP traffic is matched in the default class map.
ILS (LDAP)	TCP/389	No PAT.	
IP Options	—		All IP Options traffic is matched in the default class map.
MGCP	UDP/2427, 2727		
NETBIOS Name Server	UDP/137, 138 (Source ports)		NetBIOS is supported by performing NAT of the packets for NBNS UDP port 137 and NBDS UDP port 138.
RSH	TCP/514	No PAT.	
RTSP	TCP/554	No PAT. No outside NAT.	No handling for HTTP cloaking.
SIP	TCP/5060; UDP/5060	No outside NAT. No NAT on same security interfaces.	
Skinny Client Control Protocol (SCCP)	TCP/2000	No outside NAT. No NAT on same security interfaces.	
SMTP and ESMTP	TCP/25		
SQL*Net	TCP/1521		Versions 1 and 2.
Sun RPC over UDP	UDP/111	No NAT or PAT.	The default rule includes UDP port 111; if you want to enable Sun RPC inspection for TCP port 111, you need to create a new rule that matches TCP port 111 and performs Sun RPC inspection.
TFTP	UDP/69		Payload IP addresses are not translated.
XDMCP	UDP/177	No NAT or PAT.	

Configuring TCP Maps

Use the Add and Edit TCP Map dialog boxes to define TCP normalization maps for use with IPS, QoS, and Connection Rules service policies. The TCP normalization feature lets you specify criteria that identify abnormal packets, which the security appliance drops when they are detected. The map is used for TCP traffic that passes through the device or that is going to the device.

These TCP maps can be applied to TCP flows on PIX 7.x+ and ASA devices. For more information on configuring IPS, QoS, and Connection Rules, see Chapter 58, "Configuring Service Policy Rules on Firewall Devices".

Navigation Path

Select **Manage > Policy Objects**, then select **Maps > TCP Maps** from the Object Type selector. Right-click inside the work area and choose New Object, or right-click a row and choose Edit Object.

These dialog boxes also can be opened by clicking the Create or Edit buttons in the TCP Maps Selector while defining a Service Policy rule. See the "Connection Settings" section of Step 3. Configure the MPC actions, page 58-8 for more information about enabling TCP normalization and selecting a TCP map.

Related Topics

- Understanding Map Objects, page 6-78

Field Reference

Table 58-7 Add and Edit TCP Map Dialog Boxes

Element	Description
Name	The name of the TCP normalization map. A maximum of 128 characters is allowed.
Description	A description of the map object. A maximum of 1024 characters is allowed.
Queue Limit (ASA devices only)	The maximum number of out-of-order packets that can be buffered and put in order for a TCP connection; enter a value between 1 and 250. Enter 0 to disable this setting and use the default system queue limit, which depends on the type of traffic: • Connections for application inspection, IPS, and TCP check-retransmission have a queue limit of 3 packets. If the security appliance receives a TCP packet with a different window size, then the queue limit is dynamically changed to match the advertised setting. • For other TCP connections, out-of-order packets are passed through untouched. However, if you set the Queue Limit to 1 or higher, the number of out-of-order packets allowed for all TCP traffic matches the specified value. For application inspection, IPS, and TCP check-retransmission traffic, any advertised settings are ignored. For other TCP traffic, out-of-order packets are now buffered and put in order instead of passed through untouched.

Table 58-7 Add and Edit TCP Map Dialog Boxes (continued)

Element	Description
Time Out (ASA 7.2(4)+ devices only)	The maximum amount of time that out-of-order packets can remain in the buffer before they are dropped; enter a value between 1 and 20 seconds. The default is 4 seconds. This setting is ignored if you entered 0 for the Queue Limit.
Verify TCP Checksum	If checked, checksum verification is enabled.
Drop SYN Packets with Data	If checked, TCP SYN packets that include data are dropped.
Drop Connection on Window Variation	If checked, connections that change window size unexpectedly are dropped.
Drop Packets that Exceed Maximum Segment Size	If checked, packets that exceed the maximum segment size (MSS) set by a peer are dropped.
Check if Transmitted Data is the Same as Original	If checked, retransmit data checking is enabled.
Clear Urgent Flag	If checked, the URG (urgent) flag is cleared through the security appliance. The URG flag is used to indicate that the packet contains information that is of higher priority than other data within the stream. The TCP RFC is vague about the exact interpretation of the URG flag; therefore end systems handle urgent offsets in different ways, which may make the end system vulnerable.
Enable TTL Evasion Protection	Enables the TTL evasion protection offered by the security appliance. Do not enable this option if you want to prevent attacks that attempt to evade security policy. For example, an attacker can send a packet that passes policy with a very short TTL. When the TTL goes to zero, a router between the security appliance and the endpoint drops the packet. It is at this point that the attacker can send a malicious packet with a long TTL that appears to the security appliance to be a retransmission and is passed. To the endpoint host, however, it is the first packet that has been received. In this case, an attacker is able to succeed without security preventing the attack.
Selective Acknowledgment	
Clear Selective Ack	When checked, clears window selective acknowledgment mechanism option and allows packet. When unchecked, allows packets with single selective acknowledgment option.
Selective Ack Allow Multiple	Whether or not packets with multiple selective acknowledgment mechanism (SACK) are allowed.
Note If the selective acknowledgment options are not configured, by default, packets with a single selective acknowledgment option are allowed and packets with multiple selective acknowledgment options are dropped.	
TCP Timestamp	
Clear TCP Timestamp	When checked, clears TCP timestamp option and allows packet. When unchecked, allows packets with single TCP timestamp option. **Note** When the Clear TCP timestamp option is enabled, PAWS and RTT are disabled.

Table 58-7 Add and Edit TCP Map Dialog Boxes (continued)

Element	Description
TCP Timestamp Allow multiple	Whether or not packets with multiple TCP timestamp option are allowed.
Note If the TCP timestamp options are not configured, by default, packets with a single TCP timestamp option are allowed and packets with multiple TCP timestamp options are dropped.	
Window Scale	
Clear Window Scale	When checked, clears window scale timestamp option and allows packet. When unchecked, allows packet with single window scale option.
Window Scale Allow Multiple	Whether or not the packets with multiple window scale timestamp option are allowed.
Note If the window scale options are not configured, by default, packets with a single window scale option are allowed and packets with multiple window scale options are dropped.	
Maximum Segment Size (MSS)	
Clear MSS	When checked, clears MSS option and allows packet. When unchecked, allows packet with a single MSS option.
MSS Allow Multiple	Whether or not the packets with multiple MSS options are allowed.
Max. MSS	Enter a value for the TCP MSS limit in bytes. Valid values are between 68- 65535.
Note If the MSS options are not configured, by default, packets with a single MSS option are allowed and packets with multiple MSS options are dropped.	
Allow packets with MD5 option	Whether or not to allow packets with the MD5 option. The Allow, Allow Multiple and Clear checkboxes are available when packets with MD5 option are allowed. Allow: This allows a packet with a single MD5 option. Allow Multiple: This allows a packet with multiple MD5 options. Clear: This clears the MD5 option and allows a packet.
Note If the MD5 options are not configured, by default, packets with a single MD5 option are allowed and packets with multiple MD5 options are dropped.	
Reserved Bits	Specify how to handle TCP packets with the reserved bits set in the TCP header. The six reserved bits in the TCP header are for future use and usually have a value of 0. • Clear and Allow—Clears the reserved bits in the TCP header and allows the packet. • Allow only—Permits packets with the reserved bits set in the TCP header. • Drop—Drops packets with the reserved bits set in the TCP header.

Table 58-7 Add and Edit TCP Map Dialog Boxes (continued)

Element	Description
TCP Range Options table	The TCP Range Options table lists TCP options ranges defined for the TCP map, and the action to take for those options. The typical range numbers are 6-7, 9-18 and 20-255; the lower bound must be less than or equal to the upper bound. • To add a range, click the Add button to open the Add TCP Option Range dialog box (see Add and Edit TCP Option Range Dialog Boxes, page 58-25). • To edit a range, select it and click the Edit button. • To delete a range, select it and click the Delete button. **Note** Prior to ASA 9.6(2), the TCP values are in the range of 6-7 and 9-255.
Category	The category assigned to the map object. Categories help you organize and identify rules and objects. See Using Category Objects, page 6-13.

Add and Edit TCP Option Range Dialog Boxes

Use the Add and Edit TCP Option Range dialog boxes to define or edit a range of TCP options for use with a TCP normalization map; these are TCP options not explicitly supported on the device. This feature lets you allow or discard packets with the specified TCP options set. The typical range numbers are 6-7, 9-18 and 20-255.

Navigation Path

In the Add and Edit TCP Map dialog boxes, right-click inside the TCP Range Options table and choose Add Row, or right-click an existing row and choose Edit Row. See Configuring TCP Maps, page 58-22.

Field Reference

Table 58-8 Add and Edit TCP Option Range Dialog Boxes

Element	Description
Note Prior to ASA 9.6(2), specify either 6 or 7 or an integer from 9-255 for the lower and upper bounds of the range, respectively. For all devices, the Lower bound must be less than or equal to the Upper bound.	
Lower	The lower bound of the range; enter either 6 or 7, or an integer from 9 to 18 or an integer from 20-255.
Upper	The upper bound of the range; enter either 6 or 7, or an integer from 9 to 18 or an integer from 20-255.
Action	Choose the action to take for packets with the specified options set: • Allow – Allows any packet with a specified option set. • Clear – Clears the specified option from any packet that has it set and allows the packet. • Drop – Discards any packet with a specified option set.

Configuring Security Contexts on Firewall Devices

You can define multiple security "contexts" on a single security appliance. Each context operates as an independent virtual device, with its own security policy, interfaces and administrators. Multiple contexts are similar to having multiple stand-alone devices. Many features are supported in multiple-context mode, including routing tables, firewall features, IPS, and management. Some features are not supported; for example, VPN, multicast, and dynamic routing protocols; security contexts support only static routes; and you cannot enable OSPF or RIP in multiple-context mode. Also, some features are not directly managed by Cisco Security Manager, such as the IPS feature set for ASA and PIX devices.

In multiple-context mode, the security appliance includes a configuration for each context that identifies the security policy, interfaces, and most of the options you can configure on a stand-alone device. The system administrator adds and manages contexts by configuring them in the system configuration, which, like a single-mode configuration, is the start-up configuration. The system configuration identifies basic settings for the security appliance, but it does not include any network interfaces or network settings for itself; rather, when the system needs to access network resources (such as downloading the contexts from the server), it uses the context that is designated as the Admin context. The system configuration is used to add, delete and edit basic context settings, including allocating network interfaces to the various contexts.

The Admin context is just like any other context, except that when a user logs in to the Admin context, that user has system administrator rights and can access the system configuration and all other contexts.

This chapter contains the following topics:

- Enabling and Disabling Multiple-Context Mode, page 59-1
- Checklist for Configuring Multiple Security Contexts, page 59-3
- Managing Security Contexts, page 59-7

Enabling and Disabling Multiple-Context Mode

Cisco Security Manager does not support switching to multiple-context mode on an existing device. To perform this task, you must delete the device from Security Manager, enable multiple-context mode using a device manager or CLI input, and then add the device again to Security Manager. After the device is added in multiple-context mode, you can add, edit and delete security contexts.

Note When manually defining a multiple-context device, choose **Multi** from the Contexts list in the Operating System section of the New Device - Device Information dialog box.

Similarly, Cisco Security Manager does not support restoring an existing device to single-context mode. To perform this task, you must delete the device and any of its child contexts from Security Manager, restore single-context operation using a device manager or CLI input, and then add the device again to Security Manager.

Note When manually defining a single-context device, choose **Single** from the Contexts list in the Operating System section of the New Device - Device Information dialog box.

Related Topics

- Checklist for Configuring Multiple Security Contexts, page 59-3
- Managing Security Contexts, page 59-7
  - Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9
  - Add/Edit Security Context Dialog Box (FWSM), page 59-8

Checklist for Configuring Multiple Security Contexts

Security contexts allow a single physical device to act as multiple independent firewalls. Each security context defines a single virtual firewall, complete with its own configuration—and just as with physical devices, each security context must be correctly configured, or overall security can be compromised. Thus, defining and configuring multiple firewalls on the same physical appliance requires special care.

The following checklist outlines the basic steps necessary to configure a firewall device with multiple security contexts. Each of these steps may involve multiple substeps; all steps should be performed in the order presented. For example, you must define interfaces before configuring the various contexts.

Step	Task
Step 1	**Define interfaces and subinterfaces, or VLANs, on the physical appliance.** In this task, you define the interfaces and subinterfaces, or VLANs on FWSMs, that will be allocated to the various security contexts when you create them later. Provide physical interface parameters, such as connection type (Ethernet, GigabitEthernet, etc.), hardware Port ID, speed, and duplex mode, as well as VLAN ID if defining a subinterface. **Result**: All interfaces and subinterfaces are defined. For more information, see Configuring Firewall Device Interfaces, page 46-3.

Step	Task
Step 2	**Define an Admin context for administering the base security appliance.** This task is called out separately to ensure you define a context and IP address specifically for administration of the security appliance. The process is the same as defining a security context; however, during the process, be sure to check **Admin Context** to designate this as the administration context. In addition to being used to administer the appliance, the Admin context is used to publish syslog and SNMP messages to monitoring devices, such as the Cisco Security Monitoring, Analysis and Response System (CS-MARS), for further processing. Until you associate a specific management IP address with the Admin context, the IP address used to manage the security appliance is the one you specified when defining the device. When you specify a Management IP Address with the Admin context, it takes precedence over the one on the Device Properties page. **Result**: The Admin context is defined and associated with a physical interface. For more information, see: • Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 • Add/Edit Security Context Dialog Box (FWSM), page 59-8

Step	Task
Step 3	**Define each security context, or virtual firewall, on the base appliance.** In this task, you define individual security contexts, naming each, assigning a location for its configuration files, and allocating interfaces. Each security context represents a virtual firewall, and its definition includes the interfaces and range of associated VLAN IDs that are under its control. **Note** While the Admin context can operate as a firewall device, it is typically used as such only in single-context mode. Therefore, security contexts are treated as separate entities in this checklist. You cannot add new interfaces or modify the hardware Port value when defining a security context—you simply select previously defined interfaces for allocation to the context. **Result**: Each security context is defined and associated with a physical interface; the VLANs on which the security context will inspect traffic are also specified. For more information, see: • Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 • Add/Edit Security Context Dialog Box (FWSM), page 59-8
Step 4	**Submit/deploy to generate the virtual firewalls as children of the base appliance.** You must create the desired contexts on the security appliance before you can begin defining the individual settings of each context. To create contexts on the appliance, you must define them, and then either submit changes in Workflow mode, or deploy the changes to the security appliance in non-Workflow mode. When you create a security context, a “virtual firewall device” appears beneath the original security appliance in the Device View. Each virtual device is indicated by a related device icon with a dotted outline, and its name is the base security appliance name, underscore (_), context name. For example, the virtual device *asaMultiRouted_admin* would represent the Admin context (named “admin”) on the security appliance named “asaMultiRouted.” Similarly, *asaMultiRouted_security1* would represent the security context “security1” on the same base appliance. **Result**: Your changes are submitted or deployed (depending on the Workflow mode), which in turn creates the Admin and security contexts as children of the base security appliance. For more information, see: • Workflow and Activities Overview, page 1-20 • Submitting an Activity for Approval (Workflow Mode with Activity Approver), page 4-20 • Working with Deployment and the Configuration Archive, page 8-25
Step 5	**Define additional settings for each security context.** You can now complete the definition of each security context by selecting a virtual firewall device in the Device Selector and editing available policies, such as access rules, translation options and so on. **Result**: Each security context is fully defined, ready to operate as a virtual firewall.

Managing Security Contexts

The Security Contexts page lists security contexts configured for the selected device. You can add, edit and delete security contexts for an ASA, PIX 7.0+, or FWSM device running in multiple-context mode from this page.

Tip Deleting a security context from an FWSM device removes the security context from the running configuration of the device, but it does not delete the associated configuration file. This can cause problems if you later add another security context with the same name as the one previously deleted. This is a known issue for FWSM and is not connected to the behavior of Security Manager. A work-around is to use the CLI to delete the configuration file from the device.

Remember, the security appliance must be in multiple-context mode in order for you to configure contexts using Security Manager. See Enabling and Disabling Multiple-Context Mode, page 59-1 for more information.

Follow these steps to manage security contexts:

Step 1 Ensure Device View is your present application view; if necessary, click the **Device View** button on the toolbar.

For more information on using the Device View to configure device policies, see Managing Policies in Device View and the Site-to-Site VPN Manager, page 5-30).

Step 2 Select the appliance you want to configure.

Step 3 Select **Security Contexts** in the Device Policy selector to display the Security Contexts page.

Note The child contexts of a multiple-mode device are represented using a different icon than firewall devices in single mode.

Step 4 Add, edit and delete contexts, as necessary:

- To define a new context, click the **Add Row** button at the bottom of the page to open the Add Security Context box.
- To edit an existing context, select the desired entry in the Security Contexts list and then click the **Edit Row** button at the bottom of the page to open the Edit Security Context dialog box.
- To delete an existing context, select the desired entry in the list and then click the **Delete Row** button.

  Note Deleting a security context here will also cause the security context device to be removed from device inventory.

  Confirm the deletion of the security context and corresponding security context device.

Note Except for the titles, the Add Security Context dialog box and the Edit Security Context dialog box are identical. For PIX/ASA devices, see Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 for more information; for FWSMs, see Add/Edit Security Context Dialog Box (FWSM), page 59-8 for more information.

Add/Edit Security Context Dialog Box (FWSM)

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

The Add Security Context and Edit Security Context dialog boxes let you define and maintain contexts for the currently selected Firewall Service Module. (Except for their titles, the two dialog boxes are identical.)

Note that at least one security context must be designated as the Admin context.

Caution Security Manager does not support mapped (that is, "named" or "aliased") interfaces for FWSMs. If you discover an FWSM with named interfaces and then change the related configuration, redeployment will fail. Replace any interface aliases with the appropriate VLAN IDs.

Navigation Path

You can access the Add Security Context and Edit Security Context dialog boxes from the Security Contexts page, as described in Managing Security Contexts, page 59-7.

Field Reference

Table 59-1 Add/Edit Security Context Dialog Box (FWSM)

Element	Description
Name	Enter a name of up to 32 characters for the context. The names `System` and `Null` (in any combination of upper- and lower-case letters) are reserved, and cannot be used. **Note** While context names are case-sensitive on the device, they are not in Security Manager. That is, you cannot have two contexts with the same name but different capitalization in Security Manager.
Mode (FWSM 3.1+)	Choose the mode, Router or Transparent, for this security context. **Note** You cannot change the chosen mode in the Edit Security Context dialog box.
Admin Context	Check this box if this context is to be the Admin context for this device. **Note** The name of the Admin context for the device is displayed below the Security Contexts table.

Table 59-1 Add/Edit Security Context Dialog Box (FWSM) (continued)

Element	Description
VLAN IDs	Enter the VLANs assigned to this context. Use commas to separate multiple VLAN entries.
Config URL	Specify the context configuration location, as a URL-type address, by choosing a file-system protocol and then entering the path and name of the file to access for the context configuration. That is, choose a protocol type from the drop-down list, and then type the server name (for remote file systems), path, and file name in the related text field. For example, the combined URL for FTP has the following format: `ftp://server.example.com/configs/admin.cfg`. Available protocols are: • `disk:/` • `ftp://` • `http://` • `https://` • `tftp://`
Failover Group	If this context is part of an active/active failover configuration, choose the failover group to which this context belongs.
Description	Enter an optional description for the context.

Add/Edit Security Context Dialog Box (PIX/ASA)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

The Add Security Context and Edit Security Context dialog boxes let you define and maintain contexts for the currently selected PIX/ASA security appliance. (Except for their titles, the two dialog boxes are identical.)

Note that at least one security context must be designated as the Admin context.

Navigation Path

You can access the Add Security Context and Edit Security Context dialog boxes from the Security Contexts page, as described in Managing Security Contexts, page 59-7.

Field Reference

Table 59-2 Add/Edit Security Context Dialog Box (PIX/ASA)

Element	Description
Name	Enter a name of up to 32 characters for the context. The names `System` and `Null` (in any combination of upper- and lower-case letters) are reserved, and cannot be used. **Note** While context names are case-sensitive on the device, they are not in Security Manager. That is, you cannot have two contexts with the same name but different capitalization in Security Manager.
Description	Enter an optional description for the context.
Mode (ASA 9.0+)	Choose the mode, Router or Transparent, for this security context. **Note** You cannot change the chosen mode in the Edit Security Context dialog box.
Admin Context	Check this box if this context is to be the Admin context for this device. **Note** The name of the Admin context for the device is displayed below the Security Contexts table. **Note** If this box is checked, the IPv4 Address Pool field is disabled.
Config URL	Specify the context configuration location, as a URL-type address, by choosing a file-system protocol and then entering the path and name of the file to access for the context configuration. That is, choose a protocol type from the drop-down list, and then type the server name (for remote file systems), path, and file name in the related text field. For example, the combined URL for FTP has the following format: `ftp://server.example.com/configs/admin.cfg`. Available protocols are: • `disk0:/` • `disk1:/` • `flash:/` • `ftp://` • `http://` • `https://` • `tftp://`
VPN in multiple context mode—Beginning with Security Manager version 4.12 for ASA version 9.6(2) devices, remote access VPN on multi-context supports flash virtualization. Within a multi-context structure, each created user context can have a private storage space and a shared storage place based on the total flash that is available.	

Table 59-2 Add/Edit Security Context Dialog Box (PIX/ASA) (continued)

Element	Description
Storage URL - Private	Click the Private check box to store files associated only with that user and specific to the content that you want for that user. From the drop-down menu, choose the private directory that you created and map it to what you designated in Config URL. Select one of the following options for Private Storage URL for multi-context ASA 9.6(2) or later devices. • `disk0:/` • `flash:/` The default value of Storage URL - Private is **disk0:/**. You can modify this value. This context label name is used as a directory while performing any file deploy activity for ASA 9.6(2) Multi Context devices.
Storage URL - Shared	Click the Shared check box to upload files to the shared storage space and have it accessible to any user context for read/write access. From the drop-down menu. choose the shared directory that you created and map it to what you designated in Config URL. Select one of the following options for Shared Storage URL for multi-context ASA 9.6(2) or later devices. • `disk0:/` • `flash:/` The default value of Storage URL - Shared is **shared**. You can modify this value. This context label name is used as a directory while performing any file deploy activity for ASA 9.6(2) Multi Context devices.
ScanSafe Settings	To enable ScanSafe inspection in this context, select Enable ScanSafe Web Security. To override the license specified in the system configuration, enter a license ID in the License field; must be 32 hexadecimal characters. See Chapter 20, "Working with ScanSafe Web Security" for more information.
Interfaces	This table lists the interfaces and subinterfaces allocated to this context, and their associated settings. These are the interfaces and subinterfaces for which the security context will inspect traffic. To add interfaces and subinterfaces to this context, click the Add Row button below the table to open the Allocate Interfaces Dialog Box (PIX/ASA only), page 59-12. You can allocate one or more interfaces, and optionally with each interface, one or a range of subinterfaces. To edit an allocation entry, select it and then click the Edit Row button below the table to open the Edit Interface dialog box. Note that you can edit only the Alias Name and the Show hardware properties option; you cannot change the interface/subinterface assignments. Refer to Allocate Interfaces Dialog Box (PIX/ASA only), page 59-12 for more information about these options. To remove an interface/subinterface allocation, select the appropriate row in this table and then click the Delete Row button below the table.

Table 59-2 Add/Edit Security Context Dialog Box (PIX/ASA) (continued)

Element	Description
Failover Group	If this context is part of an active/active failover configuration, choose the failover group to which this context belongs.

Allocate Interfaces Dialog Box (PIX/ASA only)

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

The Allocate Interfaces dialog box lets you assign an interface, and optionally one or a range of related subinterfaces, to a context, and set name-aliasing options.

Navigation Path

You access the Allocate Interfaces dialog box from the Add Security Context and Edit Security Context dialog boxes. See Add/Edit Security Context Dialog Box (PIX/ASA), page 59-9 for more information.

Related Topics

- Managing Security Contexts, page 59-7

Field Reference

Table 59-3 Allocate Interfaces Dialog Box

Element	Description
Physical Interface	Choose a physical interface to assign to this context. In transparent firewall mode, you can assign only an interface that has not been allocated to another context. If you choose an interface already assigned to another context, you must also specify an subinterface.
Sub Interface ID From/To	Use these drop-down lists to specify a subinterface, or a range of subinterfaces; both lists present the subinterface IDs associated with the chosen Physical Interface. To specify a single subinterface, choose the desired ID from the first list. To specify a range, if available, choose the ending ID from the second list. (In transparent firewall mode, only subinterfaces that have not been allocated to other contexts are shown.)
View Allocation button	Click this button to open the View Interface Allocation dialog box, which presents a read-only list of all physical interfaces defined on this device and the security contexts and failover groups associated with each. You can use this to quickly determine current allocations without closing the Allocate Interfaces dialog box.

Table 59-3 Allocate Interfaces Dialog Box (continued)

Element	Description
Use aliased name in context	To enable name aliasing for this interface/subinterface, check **Use aliased names in the security context** and then enter an alias in the Alias Name field. The specified alias replaces the name of this physical interface or subinterface anywhere it would be displayed for this context—for example, the Hardware Port column on the Interfaces Page.
Alias Name	Enter the desired alias. An alias must start with a letter, end with a letter or digit, and have as interior characters only letters, digits, and underscores.
Suffix Range From/To	If you specified a range of subinterfaces, these fields are available to let you specify numeric suffixes for their aliased name. The aliased name for each subinterface consists of its sequence number from this range appended to the Alias Name you provided in the previous field. These values default to the beginning and ending subinterface ID numbers, but you can enter any valid range of numbers.
Show hardware properties in context	Select this option to allow the **show interface** CLI command to display of physical interface properties for the context even if you defined an alias. If not selected, the show interface output includes the aliased name.

CHAPTER 60

User Preferences

The User Preferences section consists of the Deployment page and the Transactional Commit page. The Deployment page provides access to the **Clear XLATE on deployment** option. The Transactional Commit page allows you to enable or disable the transactional commit model for access rules or NAT rules.

Configuring Deployment Preferences on Firewall Devices

Use the User Preferences Deployment page to specify deployment options for specific firewall devices. You can create a policy with the deployment options you want to use and then apply that policy to all devices that you want using those deployment settings.

Step 1 Do one of the following:

- (Device view) Select **Platform > User Preferences > Deployment** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > User Preferences > Deployment** from the Policy Types selector. Right-click **Deployment** and choose **New Deployment Policy** to create a policy, or select an existing policy from the Policies selector.

The Deployment page is displayed.

Step 2 Check **Clear XLATE on deployment** if you want the translation table cleared when a configuration is deployed to this device.

Select this option to send a **clear xlate** command to the firewall before changes to access lists are made. This command clears all NAT translations. By default this option is not selected.

Note This option is necessary for certain commands to take effect. If these commands are changed, you should make sure this option is enabled for the device. However, clearing the translation table disconnects all current connections that use translations.

Step 3 Click **Save** at the bottom of the page.

Configuring Transactional Commit Preferences on Firewall Devices

By default, when you change a rule-based policy (such as access rules), the changes become effective immediately. However, this immediacy comes at a slight cost in performance. The performance cost is more noticeable for very large rule lists in a high connections-per-second environment, for example, when you change a policy with 25,000 rules while the ASA is handling 18,000 connections per second.

The performance is affected because the rule engine compiles rules to enable faster rule lookup. By default, the system will also search uncompiled rules when evaluating a connection attempt so that new rules can be applied; since the rules are not compiled, the search takes longer.

Beginning with ASA 9.1(5), you can change this behavior so that the rule engine uses a transactional model when implementing rule changes, continuing to use the old rules until the new rules are compiled and ready for use. Using the transactional model, performance should not drop during the rule compilation. The following table clarifies the behavioral difference.

Model	Before Compilation	During Compilation	After Compilation
Default	Match old rules.	Match new rules. (Connections per second rate will decrease.)	Match new rules.
Transactional	Match old rules.	Match old rules. (Connections per second rate will be unaffected.)	Match new rules.

An additional benefit of the transactional model is that, when replacing an ACL on an interface, there is no gap between deleting the old ACL and applying the new one. This reduces the chances that acceptable connections will be dropped during the operation.

Tip If you enable the transactional model for a rule type, there are syslog messages to mark the beginning and the end of the compilation. These messages are numbered 780001 and following.

Step 1 Do one of the following:

- (Device view) Select **Platform > User Preferences > Transactional Commit** from the Device Policy selector.
- (Policy view) Select **PIX/ASA/FWSM Platform > User Preferences > Transactional Commit** from the Policy Types selector. Right-click **Transactional Commit** and choose **New Transactional Commit Policy** to create a policy, or select an existing policy from the Policies selector.

The Transactional Commit page is displayed.

Step 2 Enable the transactional commit model for the desired features. Options include:

- Access Group
- NAT

Step 3 Click **Save** at the bottom of the page.

PART 6

Router and Switch Device Configuration

CHAPTER 61

Managing Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Cisco Security Manager supports the management and configuration of security features and other platform-specific features on Cisco IOS access security routers. You configure these features in the form of policies, each of which defines a different aspect of the configuration of the router. For a detailed explanation of the policy paradigm used by Security Manager, see Chapter 5, "Managing Policies".

You can discover the configurations that are already defined on Cisco IOS routers. The discovery process imports the device configuration into Security Manager as policies and policy objects that you can then manage as required. For more information, see Discovering Router Policies, page 61-3.

Note Security Manager supports Cisco IOS Software Releases 12.3 and higher. However, a limited number of policies are supported for routers running Cisco IOS Software Release 12.1 or 12.2. See Configuring Routers Running IOS Software Releases 12.1 and 12.2, page 61-3.

By right-clicking a policy type in one of the policy selectors, you can assign a policy to a single router, share the policy among multiple routers, or unassign the policy from the device.

The following topics describe how to configure platform policies and interface policies on Cisco IOS routers:

- Interface polices:
  - Basic Interface Settings on Cisco IOS Routers, page 62-1
  - Advanced Interface Settings on Cisco IOS Routers, page 62-13
  - IPS Module Interface Settings on Cisco IOS Routers, page 62-22
  - CEF Interface Settings on Cisco IOS Routers, page 62-25
  - Dialer Interfaces on Cisco IOS Routers, page 62-28
  - ADSL on Cisco IOS Routers, page 62-34
  - SHDSL on Cisco IOS Routers, page 62-41
  - PVCs on Cisco IOS Routers, page 62-47
  - PPP on Cisco IOS Routers, page 62-71
- Device administration policies:
  - AAA on Cisco IOS Routers, page 63-2

  - User Accounts and Device Credentials on Cisco IOS Routers, page 63-14
  - Bridging on Cisco IOS Routers, page 63-18
  - Time Zone Settings on Cisco IOS Routers, page 63-22
  - CPU Utilization Settings on Cisco IOS Routers, page 63-25
  - HTTP and HTTPS on Cisco IOS Routers, page 63-28
  - Line Access on Cisco IOS Routers, page 63-35
  - Optional SSH Settings on Cisco IOS Routers, page 63-63
  - SNMP on Cisco IOS Routers, page 63-66
  - DNS on Cisco IOS Routers, page 63-74
  - Hostnames and Domain Names on Cisco IOS Routers, page 63-77
  - Memory Settings on Cisco IOS Routers, page 63-78
  - Secure Device Provisioning on Cisco IOS Routers, page 63-81
  - DHCP on Cisco IOS Routers, page 63-87
  - NTP on Cisco IOS Routers, page 63-96
- Identity policies:
  - 802.1x on Cisco IOS Routers, page 64-1
  - Network Admission Control on Cisco IOS Routers, page 64-8
- Logging policies:
  - Logging on Cisco IOS Routers, page 65-1
- Quality of Service:
  - Quality of Service on Cisco IOS Routers, page 66-1
- Routing policies:
  - BGP Routing on Cisco IOS Routers, page 67-1
  - EIGRP Routing on Cisco IOS Routers, page 67-8
  - OSPF Routing on Cisco IOS Routers, page 67-19
  - RIP Routing on Cisco IOS Routers, page 67-42
  - Static Routing on Cisco IOS Routers, page 67-50

Note The settings on the Policy Management page of the Security Manager Administration window determine which router platform policies can be managed with Security Manager. Any policy type that you do not select in this window does not appear on the configuration pages of Security Manager.

Configuring Routers Running IOS Software Releases 12.1 and 12.2

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Security Manager provides limited support for routers running Cisco IOS Software Releases 12.1 and 12.2 (with the exception of the ASR 1000 Series, which supports more features). You can configure the following policies on these routers:

- Access Rules (Layer 3 only). See Chapter 16, "Managing Firewall Access Rules".
- Access Rule Settings. See Chapter 16, "Managing Firewall Access Rules".
- Interfaces. See Basic Interface Settings on Cisco IOS Routers, page 62-1.
- FlexConfigs. See Chapter 7, "Managing FlexConfigs".

All other policies require Cisco IOS Software Release 12.3 or higher. For more information about supported devices, see *Supported Devices and Software Versions for Cisco Security Manager*.

Discovering Router Policies

You can discover the configurations of your Cisco IOS routers and import these configurations as policies into Security Manager. This makes it possible to add existing devices and manage them with Security Manager without having to manually configure each device policy by policy. For more information, see Adding Devices to the Device Inventory, page 3-7.

You can discover all Cisco IOS commands that can be configured with Security Manager. Discovery ignores unsupported commands, which means that they are left intact on the device even after subsequent deployments. Additionally, in cases where Security Manager can discover the command, but not all the subcommands and keywords related to that command, the unsupported elements are ignored and left intact on the device.

You can also rediscover the configurations of devices that you are already managing with Security Manager at any time. Be aware, however, that performing rediscovery overwrites the policies that you have defined in Security Manager, and is therefore not generally recommended. For more information, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Note We recommend that you perform deployment immediately after you discover the policies on a Cisco IOS router, *before* you make any changes to policies or unassign policies from the device. Otherwise, the changes that you configure in Security Manager might not be deployed to the device.

Note If a policy that is not configured in Security Manager was configured on the device using an out-of-band method (such as the CLI) between the time of the first discovery and rediscovery, we recommend that you perform deployment immediately after rediscovery.

Related Topics

- Understanding Policies, page 5-1

- Discovering Policies, page 5-12
- Working with Deployment and the Configuration Archive, page 8-25

Configuring Router Interfaces

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

This chapter contains the following topics:

- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Router Interfaces Page, page 62-7
- Advanced Interface Settings on Cisco IOS Routers, page 62-13
- Advanced Interface Settings Page, page 62-16
- IPS Module Interface Settings on Cisco IOS Routers, page 62-22
- IPS Module Interface Settings Page, page 62-23
- CEF Interface Settings on Cisco IOS Routers, page 62-25
- CEF Interface Settings Page, page 62-26
- Dialer Interfaces on Cisco IOS Routers, page 62-28
- Dialer Policy Page, page 62-31
- ADSL on Cisco IOS Routers, page 62-34
- ADSL Policy Page, page 62-37
- SHDSL on Cisco IOS Routers, page 62-41
- SHDSL Policy Page, page 62-42
- PVCs on Cisco IOS Routers, page 62-47
- PVC Policy Page, page 62-55
- PPP on Cisco IOS Routers, page 62-71
- PPP/MLP Policy Page, page 62-76

Basic Interface Settings on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

You typically add interfaces to Security Manager by performing discovery, as described in Discovering Policies, page 5-12. After you have discovered the interfaces, you can modify the properties of each interface.

You can also use Security Manager to configure physical and virtual interfaces manually. This is useful when you modify interface configurations of existing devices, and makes it possible for you to configure all the interfaces of a device before you physically add the device to the network.

Related Topics

- Available Interface Types, page 62-2
- Defining Basic Router Interface Settings, page 62-4
- Deleting a Cisco IOS Router Interface, page 62-6

Available Interface Types

Table 62-1 on page 62-2 describes the types of interfaces that can be configured on Cisco IOS routers.

Table 62-1 Router Interface Types

Type	Description
Null	Null interface.
Analysis-module	A Fast Ethernet interface that connects to the internal interface on the Network Analysis Module (NAM). **Note** You cannot configure parameters such as speed and duplex mode for this type of interface.
Async	Port line used as an asynchronous interface.
ATM	ATM interface.
BRI	ISDN BRI interface. This interface configuration propagates to each B channel. B channels cannot be configured individually. **Note** You must configure a dialer interface policy for calls to be placed on a BRI interface. For more information, see Dialer Interfaces on Cisco IOS Routers, page 62-28.
BVI	Bridge-group virtual interface. BVI interfaces are used to route traffic at Layer 3 to the interfaces in a bridge group.
Content-engine	Content engine (CE) network module interface. **Note** You cannot configure parameters such as speed and duplex mode for this type of interface. You cannot create subinterfaces for this type of interface.
Dialer	Dialer interface.
Ethernet	Ethernet IEEE 802.3 interface.
Fast Ethernet	100-Mbps Ethernet interface.
FDDI	Fiber Distributed Data Interface.
Gigabit Ethernet	1000-Mbps Ethernet interface.

Table 62-1 Router Interface Types (continued)

Type	Description
Group-Async	Main asynchronous interface. This interface type creates a single asynchronous interfaces to which other interfaces are associated. This one-to-many configuration enables you to configure all associated member interfaces by configuring the main interface.
HSSI	High-Speed Serial Interface.
Loopback	A logical interface that emulates an interface that is always up. For example, having a loopback interface on the router prevents a loss of adjacency with neighboring OSPF routers if the physical interfaces on the router go down. The name of a loopback interface must end with a number ranging from 0-2147483647. **Note** This interface type is supported on all platforms. You can create an unlimited number of loopback interfaces.
Multilink	Multilink interface. A logical interface used for multilink PPP (MLP).
Port channel	Port channel interface. This interface type enables you to bundle multiple point-to-point Fast Ethernet links into one logical link. It provides bidirectional bandwidth of up to 800 Mbps.
POS	Packet OC-3 interface on the Packet-over-SONET (POS) interface processor.
PRI	ISDN PRI interface. Includes 23/30 B-channels and one D-channel.
Serial	Serial interface.
Switch	Switch interface.
Ten Gigabit Ethernet	10000-Mbps Ethernet interface.
Token Ring	Token Ring interface.
Tunnel	Tunnel interface. **Note** You can create an unlimited number of virtual, tunnel interfaces. Valid values range from 0-2147483647.
VG-AnyLAN	100VG-AnyLAN port adapter.
VLAN	Virtual LAN subinterface.
Virtual Template	Virtual template interface. When a user dials in, a predefined configuration template is used to configure a virtual access interface; when the user is done, the virtual access interface goes down and the resources are freed for other dial-in uses.

Related Topics

- Defining Basic Router Interface Settings, page 62-4
- Deleting a Cisco IOS Router Interface, page 62-6
- Basic Interface Settings on Cisco IOS Routers, page 62-1

Defining Basic Router Interface Settings

When you define an interface or subinterface for a Cisco IOS router, you name it, specify how it is assigned an IP address, and optionally define other properties, such as the speed, maximum transmission unit (MTU), and the encapsulation type.

Note Basic interface settings are always local to the device on which they are configured. You cannot share this policy with other devices. You can, however, share advanced interface settings. For more information, see Advanced Interface Settings on Cisco IOS Routers, page 62-13.

Related Topics

- Deleting a Cisco IOS Router Interface, page 62-6

Step 1 In Device view, select **Interfaces > Interfaces** from the Policy selector.

The Router Interfaces Page, page 62-7 is displayed.

Step 2 To add a new interface or subinterface, click the Add Row button to open the Create Router Interface dialog box.

To edit an existing interface or subinterface, select it in the Interfaces table, and then click the Edit Row button to open the Edit Router Interface dialog box. Refer to Create Router Interface Dialog Box, page 62-8 for descriptions of the fields in these dialog boxes.

Step 3 Select **Enabled** to have Security Manager actively manage this interface or subinterface. If this option is deselected, the interface/subinterface definition is retained, but the interface/subinterface itself is disabled (or "shutdown").

Step 4 Choose **Interface** or **Subinterface** from the Type list.

Step 5 If you are creating an interface, enter a name for the interface. You can click **Select** to open a dialog box that will help you generate a standard name based on interface type and details about the interface's location, such as card, slot, and subinterface. For more information on using the dialog box to generate an interface name, see Interface Auto Name Generator Dialog Box, page 62-12.

Note When naming a BVI interface, use the bridge group number as the card number. Deployment will fail if you configure a BVI interface without configuring a corresponding bridge group.

Step 6 If you are creating a subinterface, provide the following:

a. **Parent**—Choose the parent interface for this subinterface.

b. **Subinterface ID**—Enter a number to identify the subinterface.

Note Security Manager configures serial subinterfaces as point-to-point, not multipoint.

Step 7 To specify a **Layer Type**, choose a Level 2 (data link) or Level 3 (network) option from this list.

Step 8 Choose a method of **IP** address assignment for this interface/subinterface, then provide additional information, as required:

- **Static IP**—Provide an **IP Address** and **Subnet Mask**.
- **DHCP**—No additional information is required.

- **PPPoE**—No additional information is required.
- **Unnumbered**—Provide the name of the interface from which an IP address is to be "borrowed."

Note Layer 2 interfaces do not support IP addresses.

Step 9 Define additional properties of the interface/subinterface:

- Use the **Negotiation** check box to enable and disable auto-negotiation for the interface.

  Auto-negotiation detects the capabilities of remote devices and negotiates the best possible performance between the two devices. When Negotiation is enabled, the Fast Ethernet Duplex and Speed options are disabled.

  Note Auto-negotiation is available only for Fast Ethernet and Gigabit Ethernet interfaces on ASR devices.

- Choose a transmission mode from the **Duplex** list. If you choose Auto, be sure the network device to which this interface is connected is set to automatically detect the transmission mode. (Auto is not available on ASRs; use auto-negotiation instead.)

Note You must configure a fixed speed to define the duplex value. Tunnel and loopback interfaces do not support this setting.

- Choose a transmission speed from the **Speed** list. If you choose Auto, be sure the network device to which this interface is connected is set to automatically detect the transmission speed. (Auto is not available on ASRs; use auto-negotiation instead.)
- Enter the maximum transmission unit (**MTU**), which defines the largest packet size, in bytes, that this interface can support.

Note Certain interface properties are set automatically, or are unavailable, depending on the interface type and the underlying port type. For example, the Speed options are available for Fast Ethernet and Gigabit Ethernet interfaces only.

Step 10 Choose an encapsulation method from the **Encapsulation** list:

- **None**—No encapsulation; no additional parameters are required.
- (Ethernet subinterfaces only) **DOT1Q**—VLAN encapsulation, as defined by the IEEE 802.1Q standard. Provide the following VLAN parameters for this subinterface:
  - Enter a VLAN ID to associate with this subinterface.

    Note All VLAN IDs must be unique among all subinterfaces configured on the same physical interface.

  - If you are defining the 802.1Q trunk interface, select Native VLAN.

Tip To configure DOT1Q encapsulation on an Ethernet interface without associating a VLAN with the subinterface, enter the **vlan-id dot1q**command using CLI commands or FlexConfigs. See Understanding FlexConfig Policies and Policy Objects, page 7-2. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the router.

- (Serial interfaces only) **Frame Relay**—IETF Frame Relay encapsulation. Provide a data-link connection identifier (DLCI) for the subinterface.

  Note Frame relay must be configured on the parent interface.

Note IETF Frame Relay encapsulation provides interoperability between a Cisco IOS router and equipment from other vendors. To configure Cisco Frame Relay encapsulation, use CLI commands or FlexConfigs.

Step 11 (Optional) Enter a description of up to 1024 characters for the interface.

Step 12 Click **OK** to save the interface/subinterface definition and close the dialog box. The new interface is displayed on the Router Interfaces page. Subinterfaces are displayed beneath the parent interface.

Deleting a Cisco IOS Router Interface

Although you can delete the definition of a virtual interface at any time, use this option with great care. If the interface is included in any policy definitions that exist for this router, deleting the interface causes these policy definitions to fail when they are deployed to the device.

Note Deleting the basic interface definition does not delete any advanced settings that are configured under **Interface > Settings > Advanced Settings**. You must delete these advanced settings separately. If you fail to do so, deployment fails.

Note Deleting the definition of a physical interface from the Router Interfaces page does not remove the interface from the device. If you perform this operation by mistake, you can perform rediscovery to restore the definition to Security Manager. For more information, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Related Topics

- Defining Basic Router Interface Settings, page 62-4
- Basic Interface Settings on Cisco IOS Routers, page 62-1

Step 1 Click the **Device View** button on the toolbar.

Step 2 Select a router from the Device selector.

Step 3 Select **Interfaces > Interfaces** from the Policy selector. The Router Interfaces page is displayed. See Table 62-2 on page 62-7 for an explanation of the fields on this page.

Step 4 Select an interface from the table, then click the **Delete** button. The interface is deleted.

Router Interfaces Page

Use the Router Interfaces page to view, create, edit, and delete interface definitions (physical and virtual) on a selected Cisco IOS router. The Router Interfaces page displays interfaces that were discovered by Security Manager as well as interfaces added manually after you added the device to the system.

Note Unlike other router policies, the Interfaces policy cannot be shared among multiple devices. The Advanced Settings policy, however, may be shared. See Local Policies vs. Shared Policies, page 5-3.

For more information, see Basic Interface Settings on Cisco IOS Routers, page 62-1.

Navigation Path

Select a Cisco IOS router from the Device selector, then select **Interfaces > Interfaces** from the Policy selector.

Related Topics

- Available Interface Types, page 62-2
- Deleting a Cisco IOS Router Interface, page 62-6
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-2 Router Interfaces Page

Element	Description
Interface Type	The interface type. Subinterfaces are displayed indented beneath their parent interface.
Interface Name	The name of the interface.
Enabled	Indicates whether the interface is currently enabled (managed by Security Manager) or disabled (shutdown state).
IP Address	The IP address of interfaces defined with a static address.
IP Address Type	The type of IP address assigned to the interface—static, DHCP, PPPoE, or unnumbered. (IP address is defined by a selected interface role.)
Interface Role	The interface roles that are assigned to the selected interface.
Add button	Opens the Create Router Interface Dialog Box, page 62-8. From here you can create an interface on the selected router.
Edit button	Opens the Create Router Interface Dialog Box, page 62-8. From here you can edit the selected interface.

Table 62-2 Router Interfaces Page (continued)

Element	Description
Delete button	Deletes the selected interfaces from the table. Ensure that the interface is not being used in any other policy before deleting it.

Create Router Interface Dialog Box

Use the Create Router Interface dialog box to create and edit physical and virtual interfaces on the selected Cisco IOS router.

Tip Interface configuration is specific to the type of device. Many of the options on this page might be greyed out for specific device or interface types because they do not apply or they are not configurable.

Navigation Path

Go to the Router Interfaces Page, page 62-7, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Deleting a Cisco IOS Router Interface, page 62-6
- Advanced Interface Settings Page, page 62-16

Field Reference

Table 62-3 Create Router Interface Dialog Box

Element	Description
Enabled	Whether the interface is enabled (no shutdown). If you deselect this option, the interface is created in the configuration but it is shut down.
Type	Specifies whether you are defining an interface or subinterface.
Name	Applies only to interfaces. The name of the interface. Enter a name manually, or click **Select** to display a dialog box for generating a name automatically. See Interface Auto Name Generator Dialog Box, page 62-12. Logical interfaces require a number after the name: • The range for dialer interfaces is 0-799. • The range for loopback interfaces is 0-2147483647. • The range for BVI interfaces is 1-255. • The only allowed value for null interfaces is 0.
Parent	Applies only to subinterfaces. The parent interface of the subinterface. Choose the parent interface from this list.

Table 62-3 Create Router Interface Dialog Box (continued)

Element	Description
Subinterface ID	Applies only to subinterfaces. The ID number of the subinterface.
IP	The method of IP address assignment for the interface: • Static IP—Defines a static IP address and subnet mask for the interface. Enter this information in the fields that appear below the option. **Note** You can define the mask using either dotted decimal (for example, 255.255.255.255) or CIDR notation (/32). See Contiguous and Discontiguous Network Masks for IPv4 Addresses, page 6-81. • DHCP—The interface obtains its IP address dynamically from a DHCP server. • PPPoE—The router automatically negotiates its own registered IP address from a central server (via PPP/IPCP). The following interface types support PPPoE: – Async – Serial – High-Speed Serial Interface (HSSI) – Dialer – BRI, PRI (ISDN) – Virtual template – Multilink • Unnumbered—The interface obtains its IP address from a different interface on the device. Choose an interface from the Interface list. This option can be used with point-to-point interfaces only. **Note** Layer 2 interfaces do not support IP addresses. Deployment fails if you define an IP address on a Layer 2 interface.
Layer Type	The OSI layer at which the interface is defined: • Unknown—The layer is unknown. • Layer 2—The data link layer, which contains the protocols that control the physical layer (Layer 1) and how data is framed before being transmitted on the medium. Layer 2 is used for bridging and switching. Layer 2 interfaces do not have IP addresses. • Layer 3—The network layer, which is primarily responsible for the routing of data in packets across logical internetwork paths. This routing is accomplished through the use of IP addresses.

Table 62-3 Create Router Interface Dialog Box (continued)

Element	Description
Negotiation	Available on ASRs; applies to Fast Ethernet and Gigabit Ethernet interfaces only. Auto-negotiation detects the capabilities of remote devices and negotiates the best possible performance between the two devices. When Negotiation is enabled, the Duplex and Speed options are disabled.
Duplex	The interface transmission mode: • None—The transmission mode is returned to its device-specific default setting. • Full—The interface transmits and receives at the same time (full duplex). • Half—The interface can transmit or receive, but not at the same time (half duplex). This is the default. • Auto—The router automatically detects and sets the appropriate transmission mode, either full or half duplex. Not available on ASRs; use auto-negotiation instead. **Note** When using Auto mode, be sure that the port on the active network device to which you connect this interface is also set to automatically negotiate the transmission mode. Otherwise, select the appropriate fixed mode. **Note** You can configure a duplex value only if you set the Speed to a fixed speed, not Auto. **Note** This setting does not apply to serial, HSSI, ATM, PRI, DSL, tunnel, or loopback interfaces.
Speed	Applies only to Fast Ethernet and Gigabit Ethernet interfaces. The speed of the interface: • None—The setting is not configurable on the device. • 10—10 megabits per second (10Base-T networks). • 100—100 megabits per second (100Base-T networks). This is the default for Fast Ethernet interfaces. • 1000—1000 megabits per second (Gigabit Ethernet networks). This is the default for Gigabit Ethernet interfaces. • Auto—The router automatically detects and sets appropriate interface speed. Not available on ASRs; use auto-negotiation. **Note** When using Auto mode, be sure that the port on the active network device to which you connect this interface is also set to automatically negotiate the transmission speed. Otherwise, select the appropriate fixed speed.

Table 62-3 Create Router Interface Dialog Box (continued)

Element	Description
MTU	The maximum transmission unit, which refers to the maximum packet size, in bytes, that this interface can handle. Valid values for serial, Ethernet, and Fast Ethernet interfaces range from 64 to 17940 bytes. Valid values for Gigabit Ethernet interfaces range from 1500 to 9216 bytes.
Encapsulation	The type of encapsulation performed by the interface: • None—No encapsulation. • DOT1Q—VLAN encapsulation, as defined by the IEEE 802.1Q standard. Applies only to Ethernet subinterfaces. • Frame Relay—IETF Frame Relay encapsulation. Applies only to serial interfaces (not serial subinterfaces). **Note** IETF Frame Relay encapsulation provides interoperability between a Cisco IOS router and equipment from other vendors. To configure Cisco Frame Relay encapsulation, use CLI commands or FlexConfigs.
VLAN ID	Applies only to subinterfaces with encapsulation type DOT1Q. The VLAN ID associated with this subinterface. The VLAN ID specifies where 802.1Q tagged packets are sent and received on this subinterface; without a VLAN ID, the subinterface cannot send or receive traffic. Valid values range from 1 to 4094. **Note** All VLAN IDs must be unique among all subinterfaces configured on the same physical interface. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. See Understanding FlexConfig Policies and Policy Objects, page 7-2. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the router.

Table 62-3 Create Router Interface Dialog Box (continued)

Element	Description
Native VLAN	Applies only when the encapsulation type is DOT1Q and you are configuring a physical interface that is meant to serve as an 802.1Q trunk interface. Trunking is a way to carry traffic from several VLANs over a point-to-point link between two devices. When selected, the Native VLAN is associated with this interface, using the ID specified in the VLAN ID field. (If no VLAN ID is specified for the Native VLAN, the default is 1.) The native VLAN is the VLAN to which all untagged VLAN packets are logically assigned by default. This includes the management traffic associated with the VLAN. If no VLAN ID is defined, the default is 1. For example, if the VLAN ID of this interface is 1, all incoming untagged packets and packets with VLAN ID 1 are received on the main interface and not on a subinterface. Packets sent from the main interface are transmitted without an 802.1Q tag. When deselected, the Native VLAN is not associated with this interface. **Note** The Native VLAN cannot be configured on a subinterface of the trunk interface. Be sure to configure the same Native VLAN value at both ends of the link; otherwise, traffic may be lost or sent to the wrong VLAN.
DLCI	Applies only to serial subinterfaces with Frame Relay encapsulation. Enter the data-link connection identifier to associate with the subinterface. Valid values range from 16 to 1007. **Note** Security Manager configures serial subinterfaces as point-to-point not multipoint.
Description	Additional information about the interface (up to 1024 characters).
Roles	The interface roles assigned to this interface. A message is displayed if no roles have yet been assigned.

Interface Auto Name Generator Dialog Box

Use the Interface Auto Name Generator dialog box to have Security Manager generate a name for the interface based on the interface type and its location in the router or switch.

Navigation Path

Go to the Create Router Interface Dialog Box, page 62-8, select **Interface** from the Type list, then click **Select** in the Name field.

Field Reference

Table 62-4 Interface Auto Name Generator Dialog Box

Element	Description
Type	The type of interface. Your selection from this list forms the first part of the generated name, as displayed in the Result field. For more information, see Available Interface Types, page 62-2.
Card	The card related to the interface. **Note** When defining a BVI interface, enter the number of the corresponding bridge group.
Slot	The slot related to the interface.
Port	The port related to the interface. **Note** The information you enter in these fields forms the remainder of the generated name, as displayed in the Result field.
Result	The name generated by Security Manager from the information you entered for the interface type and location. The name displayed in this field is read-only. **Tip** After closing this dialog box, you can edit the generated name in the Create Router Interface dialog box, if required.

Advanced Interface Settings on Cisco IOS Routers

In addition to the basic interface definitions that you can define on the Interfaces page, Security Manager provides a method for defining selected advanced settings on interfaces that support those settings.

Unlike the basic interface settings defined on the Interface page, you can share an advanced settings policy with multiple devices. This provides a convenient method for configuring multiple devices with identical settings. See Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37.

You can define a variety of advanced settings on a selected interface, subinterface, or interface role, including:

- Cisco Discovery Protocol (CDP) settings.
- Internet Control Message Protocol (ICMP) settings.
- Directed broadcast settings.
- Load interval for determining the average load.
- Throughput delay for use by routing protocols.
- Configuring TCP maximum segment size.
- Helper addresses for forwarding UDP broadcasts. For more information on helper addresses, see Understanding Helper Addresses, page 62-14.
- Enabling Maintenance Operation Protocol (MOP).
- Enabling virtual fragmentation reassembly (VFR).
- Enabling proxy ARP.
- Enabling NBAR protocol discovery.

- Enabling and configuring unicast reverse path forwarding (RFP).

Tip You can define these settings for multiple interfaces on a device at once by choosing an interface role instead of a specific interface. For example, if you have defined an All-Ethernets interface role, you can define identical advanced settings for every Ethernet interface on the device with a single definition. See Understanding Interface Role Objects, page 6-73.

Before You Begin

- Define basic interface settings. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > Advanced Settings** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > Advanced Settings** from the Policy Type selector. Select an existing policy or create a new one.

The Advanced Interface Settings page is displayed (see Advanced Interface Settings Page, page 62-16).

Step 2 Do one of the following:

- Click the **Add** button to add an interface or interface role to the table. In the Advanced Interface Settings dialog box, enter the name of the interface or interface role, or click **Select** to select an existing role or to create a new role.
- Select an existing entry in the table and click the **Edit** button to change its settings.

Step 3 Configure the advanced settings required for the selected interface. For details about each setting, see Advanced Interface Settings Dialog Box, page 62-16.

Step 4 Click **OK** to save your definitions. Your definitions are displayed in the Advanced Interface Settings table.

Understanding Helper Addresses

Network hosts occasionally use User Datagram Protocol (UDP) broadcasts to determine address, configuration, and name information. This presents a problem if the host is on a network segment that does not include the required server, as by default, routers do not forward UDP broadcasts beyond their subnet. You can remedy this situation by configuring the interface to forward certain classes of broadcasts to a helper address.

One common use of helper addresses is when the router acts as a relay agent for DHCP clients who need to contact a DHCP server located on a different subnet. The helper address can either represent a specific DHCP server or a network address for a segment containing multiple DHCP servers. You can also configure a helper address for each DHCP server.

In Figure 62-1, hosts located on network 192.168.1.0 can use 10.44.23.7 as a helper address to forward UDP broadcasts to the other network, while hosts located on network 10.44.0.0 can use 192.168.1.19 as their helper address.

Figure 62-1 Helper Addresses

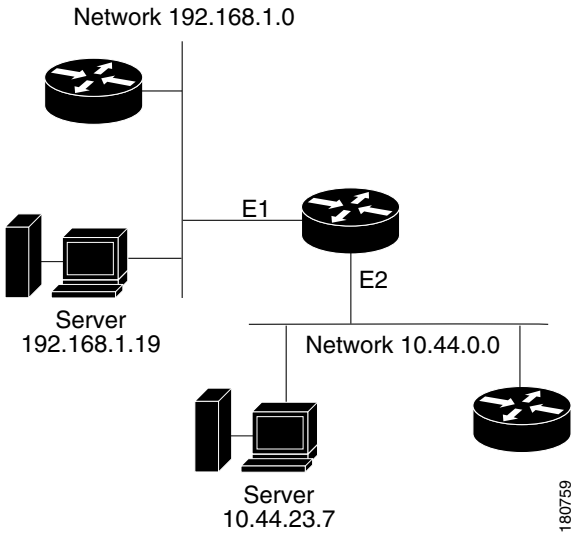

Table 62-5 on page 62-15 lists the default UDP services that can be forwarded to helper addresses.

Table 62-5 Default UDP Services Forwarded to Helper Addresses

Service	Port
BOOTP/DHCP Client	68
BOOTP/DHCP Server	67
DNS	53
NetBIOS datagram service	138
NetBIOS name service	137
TACACS	49
TFTP	69
Time	37

Tip To forward additional UDP services, use the CLI or FlexConfigs to configure the ip forward-protocol command. Use the no form of this command to prevent the forwarding of any of the default services listed in Table 62-5 on page 62-15.

All of the following conditions must be met in order for a UDP or IP packet to use helper addresses:

- The MAC address of the received frame must be an all-ones broadcast address (ffff.ffff.ffff).
- The IP destination address must be one of the following: all-ones broadcast (255.255.255.255), subnet broadcast for the receiving interface, or major-net broadcast for the receiving interface if the no ip classless command is also configured.
- The IP time-to-live (TTL) value must be at least 2.
- The IP protocol must be UDP (17).

Related Topics

- Advanced Interface Settings on Cisco IOS Routers, page 62-13
- Basic Interface Settings on Cisco IOS Routers, page 62-1

Advanced Interface Settings Page

Use the Advanced Interface Settings page to configure advanced interface definitions (physical and virtual) on a router. Examples of advanced settings include Cisco Discovery Protocol (CDP) settings, ICMP message settings, and virtual fragment reassembly settings. You can configure settings for specific interfaces or for interface roles. The columns in the table summarize the advanced settings for an entry and are explained in Advanced Interface Settings Dialog Box, page 62-16.

To configure advanced settings:

- Click the **Add** button to add an interface or interface role to the table, and fill in the Advanced Interface Settings dialog box.
- Select an entry and click the **Edit** button to edit an existing entry.
- Select an entry and click the **Delete** button to delete it.

For more information, see Advanced Interface Settings on Cisco IOS Routers, page 62-13.

Navigation Path

- (Device view) Select **Interfaces > Settings > Advanced Settings** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > Advanced Settings** from the Policy Type selector. Right-click **Advanced Settings** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Router Interfaces Page, page 62-7
- Available Interface Types, page 62-2
- Deleting a Cisco IOS Router Interface, page 62-6
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Advanced Interface Settings Dialog Box

Use the Advanced Interface Settings dialog box to define a variety of advanced settings on a selected interface as described in the table below.

Navigation Path

Go to the Advanced Interface Settings Page, page 62-16, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Advanced Interface Settings on Cisco IOS Routers, page 62-13

- Deleting a Cisco IOS Router Interface, page 62-6
- Available Interface Types, page 62-2

Field Reference

Table 62-6 Advanced Interface Settings Dialog Box

Element	Description
Interface	The interface on which the advanced settings are defined. Enter the name of an interface or interface role, or click **Select** to select it. If the you want is not listed, click the **Create** button to create it. **Note** The only advanced settings supported on Layer 2 interfaces are Max. Bandwidth, Load Interval, and CDP.
Max Bandwidth	The bandwidth value to communicate to higher-level protocols in kilobits per second (kbps). The value you define in this field is an informational parameter only; it does not affect the physical interface.
Load Interval	The length of time, in seconds, used to calculate the average load on the interface. Valid values range from 30 to 600 seconds, in multiples of 30 seconds. The default is 300 seconds (5 minutes). Load interval is not supported on subinterfaces. Modify the default to shorten the length of time over which load averages are computed. You can do this if you want load computations to be more reactive to short bursts of traffic. Load data is gathered every 5 seconds. This data is used to compute load statistics, including input/output rate in bits and packets per second, load, and reliability. Load data is computed using a weighted-average calculation in which recent load data has more weight in the computation than older load data. **Tip** You can use this option to increase or decrease the likelihood of activating a backup interface; for example, a backup dial interface may be triggered by a sudden spike in the load on an active interface.
TCP Maximum Segment Size	The maximum segment size (MSS) of TCP SYN packets that pass through this interface. Valid values range from 500 to 1460 bytes. If you do not specify a value, the MSS is determined by the originating host. This option helps prevent TCP sessions from being dropped as they pass through the router. Use this option when the ICMP messages that perform auto-negotiation of TCP frame size are blocked (for example, by a firewall). We highly recommend using this option on the tunnel interfaces of DMVPN networks. For more information, see *TCP MSS Adjustment* at this URL: http://www.cisco.com/en/US/docs/ios/12_2t/12_2t4/feature/guide/ft_admss.html **Note** Typically, the optimum MSS is 1452 bytes. This value plus the 20-byte IP header, the 20-byte TCP header, and the 8-byte PPPoE header add up to a 1500-byte packet that matches the MTU size for the Ethernet link.

Table 62-6 Advanced Interface Settings Dialog Box (continued)

Element	Description
Helper Addresses	The helper addresses that are used to forward User Datagram Protocol (UDP) broadcasts that are received on this interface. Enter one or more addresses or the names of the network/host objects, or click **Select** to select an object from a list or to create a new object. By default, routers do not forward broadcasts outside of their subnet. Helper addresses provide a solution by enabling the router to forward certain types of UDP broadcasts as a unicast to an address on the destination subnet. For more information, see Understanding Helper Addresses, page 62-14.
Interface Throughput Delay	The expected delay for the interface in tens of microseconds (for example, 3000 translates to 30,000 microseconds). You can enter a value between 1 and 16777215, and the default varies by the type of interface. Higher-level protocols might use delay information to make operating decisions. For example, IGRP can use delay information to differentiate between a satellite link and a land link. This setting is for informational purposes only and does not affect the actual delay on the interface.
Cisco Discovery Protocol settings	Settings related to the Cisco Discovery Protocol (CDP). CDP is a media- and protocol-independent device-discovery protocol that runs on all Cisco-manufactured equipment including routers, access servers, bridges, and switches. It is primarily used to obtain protocol addresses of neighboring devices and to discover the platform of those devices. The options are: • Enable CDP—Whether to enable the Cisco Discovery Protocol (CDP) on this interface. You cannot enable CDP on ATM interfaces. • Log CDP Messages—On Ethernet interfaces, whether to log duplex mismatches for this interface.
ICMP Messages Settings	
Enable Redirect Messages	Whether to enable the sending of Internet Control Message Protocol (ICMP) redirect messages if the device is forced to resend a packet through the same interface on which it was received to another device on the same subnet. Redirect messages are sent when the device wants to instruct the originator of the packet to remove it from the route and substitute a different device that offers a more direct path to the destination.

Table 62-6 Advanced Interface Settings Dialog Box (continued)

Element	Description
Enable Unreachable Messages	Whether to enable the sending of ICMP unreachable messages. Unreachable messages are sent in two circumstances: • If the interface receives a nonbroadcast packet destined for itself that uses an unknown protocol, the interface sends an ICMP unreachable message to the source. • If the device receives a packet that it cannot deliver to its ultimate destination because it knows of no route to the destination address, it sends an ICMP host unreachable message to the originator of the packet. **Note** This is the only advanced setting supported by the null0 interface.
Enable Mask Reply Messages	Whether to enable the sending of ICMP mask reply messages. Mask reply messages are sent in response to mask request messages, which are sent when a device needs to know the subnet mask for a particular subnetwork.
Additional Settings	
Enable Maintenance Operation Protocol (MOP)	Whether to enable MOP on the interface. You can use MOP for utility services such as uploading and downloading system software, remote testing, and problem diagnosis.
Enable Virtual Fragment Reassembly (VFR)	Whether to enable virtual fragmentation reassembly (VFR) on this interface. VFR is a feature that enables the Cisco IOS Firewall to create dynamic ACLs that can protect the network from various fragmentation attacks. For more information, see *Virtual Fragmentation Reassembly* at this URL: http://www.cisco.com/en/US/docs/ios/sec_data_plane/configuration/guide/sec_virt_frag_reassm_ps6441_TSD_Products_Configuration_Guide_Chapter.html
Enable Proxy ARP	Whether to enable proxy Address Resolution Protocol (ARP) on the interface. Proxy ARP, defined in RFC 1027, is the technique in which one host, usually a router, answers ARP requests intended for another machine, thereby accepting responsibility for routing packets to the real destination. Proxy ARP can help machines on a subnet reach remote subnets without configuring routing or a default gateway.
Enable NBAR Protocol Discovery	Whether to enable network-based application recognition (NBAR) on this interface to discover traffic and keep traffic statistics for all protocols known to NBAR. Protocol discovery provides a method to discover application protocols traversing an interface so that QoS policies can be developed and applied to them. For more information, go to: http://www.cisco.com/en/US/products/ps6616/products_qanda_item09186a00800a3ded.shtml

Table 62-6 Advanced Interface Settings Dialog Box (continued)

Element	Description
Enable Directed Broadcasts ACL	Whether to have directed broadcast packets "exploded" as a link-layer broadcast when this interface is directly connected to the destination subnet. When deselected, directed broadcast packets that are intended for the subnet to which this interface is directly connected are dropped rather than being broadcast. This is the default. An IP directed broadcast is an IP packet whose destination address is a valid broadcast address on a different subnet from the node on which it originated. In such cases, the packet is forwarded as if it was a unicast packet until it reaches its destination subnet. This option affects only the final transmission of the directed broadcast on its destination subnet; it does not affect the transit unicast routing of IP directed broadcasts. If you enable directed broadcasts, you can apply an ACL to determine which directed broadcasts are permitted to be broadcast on the destination subnet. All other directed broadcasts destined for the subnet to which this interface is directly connected are dropped. Enter the name of a standard or extended ACL object, or click **Select** to select an object from a list or to create a new object. **Tip** Because directed broadcasts, and particularly ICMP directed broadcasts, have been abused by malicious persons, we recommend deselecting this option on interfaces where directed broadcasts are not needed. When you enable directed broadcasts, apply an ACL to restrict their use.
Unicast Reverse Path Forwarding (RFP) Settings	
Enable Unicast RFP	Whether to enable unicast reverse path forwarding (RFP) on the interface. When you enable Unicast RPF on an interface, the router examines all packets that are received on that interface. The router checks to make sure that the source address appears in the FIB, and takes action based on your unicast RFP settings. Use unicast RFP to mitigate problems caused by malformed or forged (spoofed) IP source addresses that pass through a router. Malformed or forged source addresses can indicate DoS attacks based on source IP address spoofing. For more information on unicast RFP, see the description of the **ip verify unicast source reachable-via** command in the *Cisco IOS Interface and Hardware Component Command Reference*. To enable unicast RFP, you must also globally enable Cisco Express Forwarding (CEF). For more information on CEF, see CEF Interface Settings on Cisco IOS Routers, page 62-25.

Table 62-6 Advanced Interface Settings Dialog Box (continued)

Element	Description
Mode	How strict to make unicast RFP: • Loose Mode—The default. Examines incoming packets to determine whether the source address is in the Forwarding Information Base (FIB) and permits the packet if the source is reachable through any interface on the router. Use loose mode on interfaces where asymmetric paths allow packets from valid source networks (networks contained in the FIB). For example, routers that are in the core of an ISP network have no guarantee that the best forwarding path out of the router will be the path selected for packets returning to the router. • Strict Mode—Examines incoming packets to determine whether the source address is in the FIB and permits the packet only if the source is reachable through the interface on which the packet was received. Use strict mode on interfaces where only one path allows packets from valid source networks (networks contained in the FIB). Also, use strict mode when a router has multiple paths to a given network as long as the valid networks are switched through the incoming interfaces. Packets for invalid networks are dropped. For example, routers at the edge of the network of an ISP are likely to have symmetrical reverse paths. Strict mode is also applicable in certain multihomed situations, provided that optional Border Gateway Protocol (BGP) attributes, such as weight and local preference, are used to achieve symmetric routing.
Allow Use Of Default Route for RFP Verification	Whether to permit Unicast RPF to successfully match on prefixes that are known through the default route when determining whether to pass packets. Normally, sources found in the FIB but only by way of the default route are dropped.
Allow Self Ping	Whether to allow the router to ping its own interfaces. By default, when you enable Unicast RPF, packets that are generated by the router and destined to the router are dropped, thereby making certain troubleshooting and management tasks difficult to accomplish. ⚠ **Caution** Allowing self-ping opens a potential denial of service (DoS) hole.
ACL (For Unicast RFP)	If you enable unicast RFP, you can apply an ACL to refine how packets are handled when a reverse path is not found. If you specify an ACL, when (and only when) a packet fails the Unicast RPF check, the ACL is checked to determine whether the packet should be dropped (using a deny statement in the ACL) or forwarded (using a permit statement in the ACL). Enter the name of a standard or extended ACL object, or click **Select** to select an object from a list or to create a new object.

IPS Module Interface Settings on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IPS and IOS features/functionality, it does not support any bug fixes or enhancements.

On some routers, you can install IPS modules such as the Cisco Intrusion Prevention System Advanced Integration Module or Network Module. When installed and active, you must configure the IPS Module interface settings policy to define the following:

- The name of the interface between the module and the router.
- The failure mode of the module. If the module fails, you can configure it to allow all traffic or to deny all traffic.
- The router interfaces to monitor. You can name specific interfaces or use interface roles to cover more than one interface at a time. For example, if you have defined an All-Ethernets interface role, you can define identical monitoring settings for every Ethernet interface on the device with a single definition. See Understanding Interface Role Objects, page 6-73.

Tip After you have defined an IPS Module interface settings policy, you can share the policy and assign it to other devices. This provides a convenient method for configuring multiple devices with identical settings. See Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37.

Before You Begin

Define basic interface settings. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > IPS Module** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > IPS Module** from the Policy Type selector. Select an existing policy or create a new one.

The IPS Module Interface Settings page is displayed. See IPS Module Interface Settings Page, page 62-23 for an explanation of the fields on this page.

Step 2 In the IPS Module Interface Settings fields, enter the name of the IPS interface (such as IDS-Sensor1/0) or click Select to select it from a list. Also determine whether you want to allow all traffic if the module fails (fail open) or to deny all traffic (fail closed).

Step 3 Identify the router interfaces that the module should monitor. Click the **Add** button below the IPS Module Service Module Monitoring Settings table to add interfaces to the list, or select an interface and click the **Edit** button to change the settings for an existing interface. Use the IPS Monitoring Information dialog box to define the interface name or role, monitoring mode, and access list (if any). For more information, see IPS Monitoring Information Dialog Box, page 62-24.

IPS Module Interface Settings Page

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Use the IPS Module Interface Settings page to define the settings on the Cisco Intrusion Prevention System Advanced Integration Module or Network Module. The module must be running IPS 6.0 or later. You can define the fail mode for the IPS interface, and the interfaces that the module should monitor. Configure this policy only if the router hosts an IPS module.

Caution Cisco IOS IPS and the Cisco IPS module cannot be used together. Cisco IOS IPS must be disabled when the IPS module is installed.

Navigation Path

- (Device view) Select **Interfaces > Settings > IPS Module** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > IPS Module** from the Policy Type selector. Create a new policy or select an existing policy from the Shared Policy selector.

Related Topics

- IPS Module Interface Settings on Cisco IOS Routers, page 62-22
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-7 IPS Module Interface Settings Page

Element	Description
Interface Name	The name of the IPS module interface. Enter the name or click **Select** to select the interface or interface role. If the object that you want is not listed, click the **Create** button to create it.
Fail Over Mode	How the module should handle traffic inspection during a module failure, either to fail open (passing all traffic without inspection) or fail closed (dropping all traffic). The default is fail open.

Table 62-7 IPS Module Interface Settings Page (continued)

Element	Description
IPS Module Service Module Monitoring Settings table	The list of interfaces on the router that the IPS module should monitor. The table shows the name of the interface or interface role, whether monitoring is inline or promiscuous, and whether an ACL is used to filter traffic for inspection on the interface. Inline mode puts the IPS module directly into the traffic flow, allowing it to stop attacks by dropping malicious traffic before it reaches the intended target. In promiscuous mode, packets do not flow through the sensor; the sensor analyzes a copy of the monitored traffic rather than the actual forwarded packet. If the ACL is matched, the matched traffic is not inspected. • To add an interface to the table, click the Add button and fill in the IPS Monitoring Information Dialog Box, page 62-24. • To edit the settings for an interface, select it and click the Edit button. • To delete an interface, select it and click the Delete button.

IPS Monitoring Information Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Use the IPS Monitoring Information dialog box to add or edit the properties of interfaces to be monitored by the IPS module.

Navigation Path

Go to the IPS Module Interface Settings Page, page 62-23, then click the **Add** or **Edit** button beneath the IPS Module Service Module Monitoring Settings table.

Related Topics

- IPS Module Interface Settings on Cisco IOS Routers, page 62-22
- Basic Interface Settings on Cisco IOS Routers, page 62-1

Field Reference

Table 62-8 IPS Monitoring Information Dialog Box

Element	Description
Interface Name	A name of the interface or interface role that the module should monitor. Enter the name or click **Select** to select the interface or interface role. If the object that you want is not listed, click the **Create** button to create it.

Table 62-8 IPS Monitoring Information Dialog Box (continued)

Element	Description
Monitoring Mode	How the interface should be monitored: • Inline mode—The IPS module is directly in the traffic flow, allowing it to stop attacks by dropping malicious traffic before it reaches the intended target. • Promiscuous mode—Packets do not flow through the sensor; the sensor analyzes a copy of the monitored traffic rather than the actual forwarded packet.
Access List	The name of the standard or extended access list policy object to use to filter traffic on this interface for inspection, if you want to apply one. A matched ACL causes traffic not to be inspected for that ACL. Click **Select** to select the ACL or to create a new one.

CEF Interface Settings on Cisco IOS Routers

Cisco Express Forwarding (CEF) is an advanced Layer 3 IP switching technology that optimizes network performance and scalability for all kinds of networks, from those that carry small amounts of traffic to those that carry large amounts of traffic in complex patterns, such as the Internet and networks characterized by intensive web-based applications or interactive sessions. CEF is enabled by default on most Cisco IOS routers.

Typically, you do not need to configure a CEF policy unless you want to enable CEF accounting so that you can view statistics with the **show ip cef** command on the router. You would also configure the policy if you want to disable CEF, or to configure non-default CEF behavior on specific interfaces, for example, to have CEF load balance based on packets rather than source-destination packet streams.

When configuring alternate CEF settings for interfaces, you can name specific interfaces or use interface roles to cover more than one interface at a time. For example, if you have defined an All-Ethernets interface role, you can define identical CEF settings for every Ethernet interface on the device with a single definition. See Understanding Interface Role Objects, page 6-73.

Tip After you have defined a CEF interface settings policy, you can share the policy and assign it to other devices. This provides a convenient method for configuring multiple devices with identical settings. See Working with Shared Policies in Device View or the Site-to-Site VPN Manager, page 5-37.

Before You Begin

Define basic interface settings. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > CEF** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > CEF** from the Policy Type selector. Select an existing policy or create a new one.

The CEF Interface Settings page is displayed. See CEF Interface Settings Page, page 62-26 for an explanation of the fields on this page.

Step 2 If you are enabling CEF, select the accounting options you desire.

Step 3 If you want to configure non-default behavior for certain interfaces, add them to the CEF Interface Settings table. Click the **Add** button below the table to add interfaces to the list, or select an interface and click the **Edit** button to change the settings for an existing interface. For more information about the options, see CEF Interface Settings Dialog Box, page 62-27.

CEF Interface Settings Page

Use the CEF Interface Settings page to define the settings for Cisco Express Forwarding. CEF is an advanced Layer 3 IP switching technology that optimizes network performance and scalability for all kinds of networks, from those that carry small amounts of traffic to those that carry large amounts of traffic in complex patterns, such as the Internet and networks characterized by intensive web-based applications or interactive sessions. CEF is enabled by default on most Cisco IOS routers.

Navigation Path

- (Device view) Select **Interfaces > Settings > CEF** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > CEF** from the Policy Type selector. Create a new policy or select an existing policy from the Shared Policy selector.

Related Topics

- CEF Interface Settings on Cisco IOS Routers, page 62-25
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-9 CEF Interface Settings Page

Element	Description
Enable Cisco Express Forwarding	Whether to enable CEF globally on the device. The option is greyed out if you cannot disable CEF on the device. You can configure other settings on the page only if you enable CEF globally.

Table 62-9 CEF Interface Settings Page (continued)

Element	Description
CEF Network Accounting	These options are for configuring CEF accounting globally. If you collect accounting statistics, you can view them using the **show ip cef** command on the router. You can select the following options to enable different types of accounting: • Enable Accounting for Traffic Through Non-Recursive Prefixes—For network prefixes with directly connected next hops, non-recursive accounting enables express forwarding of the collection of packets through a prefix. • Enable Per-Prefix Accounting—Accounting statistics based on the packet's network prefix. • Enable Prefix Length Accounting—Accounting statistics based on the network prefix length. • Enable Load Balance Hash Accounting—When you use per-destination load balancing (the default), CEF uses a series of 16 hash buckets to distribute the available paths based on the source and destination addresses. Enabling load balance hash accounting provides per-hash-bucket counters.
CEF Interface Settings table	The interfaces on the router for which you are defining special CEF configurations. When you enable CEF globally, by default, all interfaces on the router enable CEF and use per-destination load balancing. Add interfaces to this table only if you want to configure different behavior for the interfaces. The table shows the name of the interface or interface role, whether CEF is enabled or disabled, and whether the interface is load balancing based on destination or on a per-packet basis. For a detailed explanation of the fields, see CEF Interface Settings Dialog Box, page 62-27. • To add an interface to the table, click the Add button. • To edit the settings for an interface, select it and click the Edit button. • To delete an interface, select it and click the Delete button.

CEF Interface Settings Dialog Box

Use the CEF Interface Settings dialog box to add or edit the CEF properties of interfaces when you want to configure something different than the global default.

Navigation Path

Go to the CEF Interface Settings Page, page 62-26, then click the **Add** or **Edit** button beneath the CEF Interface Settings table.

Related Topics

- CEF Interface Settings on Cisco IOS Routers, page 62-25
- Basic Interface Settings on Cisco IOS Routers, page 62-1

Field Reference

Table 62-10 CEF Interface Settings Dialog Box

Element	Description
Interface Name	The name of the interface or interface role for which you are configuring CEF. Enter the name or click **Select** to select the interface or interface role. If the object that you want is not listed, click the **Create** button to create it.
Enable CEF on Interface	Whether to enable CEF on the interface. CEF is enabled by default.
Load Balancing	How the interface should balance traffic, either per-destination or per-packet. In per-destination load balancing, all packets for a given source-destination pair take the same path. In per-packet load balancing, packets for a given source-destination pair can take different equal-cost routes, and thus reach their destination out of order. The default is to balance the load based on the destination of the traffic.

Dialer Interfaces on Cisco IOS Routers

Before you can configure a dial backup policy for a site-to-site VPN (see Configuring Dial Backup, page 25-40), you must configure a dialer interface policy on the appropriate Cisco IOS router. The dialer interface policy uses dialer pools to associate the dialer interface used by dial backup with a physical BRI interface on the router. Each dialer interface is associated with a single dialer pool, which can contain one or more physical interfaces. Multiple dialer interfaces can reference the same dialer pool.

The following topics describe how to create dialer interfaces policies on Cisco IOS routers:

- Defining Dialer Profiles, page 62-28
- Defining BRI Interface Properties, page 62-30

Defining Dialer Profiles

When you configure a dialer profile, you must select the interface or interface role representing the dialer interface and specify the number to be dialed. You must also assign a pool ID, which you use to reference this dialer interface when configuring the physical dialer interface. Additionally, you can modify the default timeout settings for the line.

Note IP is the only protocol supported for dialer profiles by Security Manager.

Note Authentication parameters for the dialer profile are defined in the PPP policy.

Before You Begin

Define the virtual and physical dialer interfaces on the router. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Note In addition, you can optionally define interface roles for the virtual and physical dialer interfaces. See Defining Dialer Profiles, page 62-28.

Related Topics

- Defining BRI Interface Properties, page 62-30
- Dialer Interfaces on Cisco IOS Routers, page 62-28

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > Dialer** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > Dialer** from the Policy Type selector. Select an existing policy or create a new one.

The Dialer page is displayed. See Table 62-11 on page 62-31 for a description of the fields on this page.

Step 2 Select a dialer profile from the upper table on the Dialer Interfaces page, then click **Edit**, or click **Add** to create a profile. The Dialer Profile dialog box appears. See Table 62-12 on page 62-32 for a description of the fields in this dialog box.

Step 3 Enter the name of the interface or interface role that represents the virtual dialer interface, or click **Select** to select an interface role object or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 Enter a name for the dialer profile. Having a name makes it easier for you to assign the correct dialer pool to the physical interface. See Defining BRI Interface Properties, page 62-30.

Tip We recommend that you define a name that is logically associated with the site to which the dialer interface serves as a backup. For example, if the dialer interface is serving as a backup connection to the London site, define the name London for the dialer profile.

Step 5 Enter an ID number for the dialer pool to associate with this dialer interface. Each dialer interface is associated with a single pool. Multiple interfaces may, however, be associated with the same dialer pool.

Step 6 Enter the number of the dialer group to assign to the dialer interface.

Step 7 (Optional) In the Interesting Traffic ACL field, enter the name of the extended ACL object that defines which packets are permitted to initiate calls using this dialer profile, or click **Select** to select the object from a list or to create a new one. Use this option to limit the IP traffic that can make use of the dialer.

Step 8 Enter the dialer string, which is the phone number of the remote side of the dialer interface connection.

Step 9 (Optional) Modify the default timeout values (Idle Timeout and Fast Idle Timeout), if required.

Step 10 Click **OK** to save your definitions locally on the client and close the dialog box. The dialer profile appears in the Dialer Profile table on the Dialer page.

Defining BRI Interface Properties

You configure the properties of the physical BRI interfaces used for dialer interface policies by selecting the appropriate interface or interface role, defining the dialer pools to which the interface belongs, and defining the ISDN switch type. It is the dialer pool that connects the physical interface with the virtual dialer interface.

Note To define other types of physical dialer interfaces, such as ATM and Ethernet, use FlexConfigs. For more information, see Understanding FlexConfig Policies and Policy Objects, page 7-2.

Before You Begin

Define the virtual and physical dialer interfaces on the router. SeeBasic Interface Settings on Cisco IOS Routers, page 62-1.

Note In addition, you can optionally define interface roles for the virtual and physical dialer interfaces. See Creating Interface Role Objects, page 6-74.

Related Topics

- Defining Dialer Profiles, page 62-28
- Dialer Interfaces on Cisco IOS Routers, page 62-28

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > Dialer** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > Dialer** from the Policy Type selector. Select an existing policy or create a new one.

The Dialer Interfaces page is displayed. See Table 62-11 on page 62-31 for a description of the fields on this page.

Step 2 Select a physical BRI interface from the Dialer Physical Interfaces table, then click **Edit**, or click **Add** to add an interface. The Dialer Physical Interface dialog box appears. See Table 62-13 on page 62-33 for a description of the fields in this dialog box.

Step 3 Enter the name of the interface or interface role that represents the physical dialer interface, or click **Select** to select an interface role object from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 Enter the names of the dialer pools to associate with the physical interface, or click **Select** to display a selector. Separate multiple entries with commas.

Step 5 Select the ISDN switch type used by the physical interface. Table 62-13 on page 62-33 describes the available switch types.

Step 6 (Optional) If you selected the Basic-DMS-100, Basic-NI, or Basic-5ess switch type, enter up to two service provider identifiers (SPIDs).

Note We recommend that you do not enter SPIDs for the Basic-5ess switch type, even though SPIDs are supported.

Step 7 Click **OK** to save your definitions locally on the client and close the dialog box. The interface definition appears in the Dialer Physical Interfaces table on the Dialer Interface page.

Dialer Policy Page

Use the Dialer page to define the relationship between physical Basic Rate Interface (BRI) and virtual dialer interfaces. You use these dialer interfaces when you configure the dial backup feature for site-to-site VPNs.

For more information, see Dialer Interfaces on Cisco IOS Routers, page 62-28.

Navigation Path

- (Device view) Select **Interfaces > Settings > Dialer** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > Dialer** from the Policy Type selector. Right-click **Dialer** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Configuring Dial Backup, page 25-40
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-11 Dialer Page

Element	Description
Dialer Profiles table	The dialer profiles that define the dialer pools. You must add profiles before you can add physical BRI interfaces. The table shows the name of the interface or interface role that the dialer interface uses, the profile name, pool, group, the ACL that defines which traffic can use this profile, the dial string, and idle times. • To add a profile, click the **Add Row** button and fill in the Dialer Profile Dialog Box, page 62-32. • To edit a profile, select it and click the **Edit Row** button. • To delete a profile, select it and click the **Delete Row** button.
Dialer Physical Interfaces (BRI) table	The physical interfaces that use the dialer profiles. The table shows the name of the interface or interface role, the dial pools, ISDN switch type, and first and second service provider identifiers (SPID) related to the interface. • To add an interface, click the **Add Row** button and fill in the Dialer Physical Interface Dialog Box, page 62-33. • To edit an interface, select it and click the **Edit Row** button. • To delete an interface, select it and click the **Delete Row** button.

Dialer Profile Dialog Box

Use the Dialer Profile dialog box to add or edit dialer profiles.

Navigation Path

Go to the Dialer Policy Page, page 62-31, then click the **Add** or **Edit** button beneath the Dialer Profile table.

Related Topics

- Dialer Physical Interface Dialog Box, page 62-33
- Defining Dialer Profiles, page 62-28
- Dialer Interfaces on Cisco IOS Routers, page 62-28
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Creating Interface Role Objects, page 6-74

Field Reference

Table 62-12 Dialer Profile Dialog Box

Element	Description
Name	A descriptive name for the dialer profile. This name enables you to assign the correct dialer pool to the physical interface. You can also use the profile name as a reference to the site to which this dialer interface serves as a backup.
Interface	The virtual dialer interface to associate with the dialer profile. Enter the name of an interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Pool ID	The dialer pool ID. Each pool can contain multiple physical interfaces and can be associated with multiple dialer interfaces. Each dialer interface, however, is associated with only one pool.
Group	The group ID, which identifies the dialer group that this dialer interface uses.
Interesting Traffic ACL	The extended, numbered ACL that defines which packets are permitted to initiate calls using this dialer profile. The valid ACL number range is 100 to 199. Enter the name of the ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Dialer String (Remote Phone Number)	The phone number of the destination that the dialer contacts.
Idle Timeout	The default amount of idle time before an uncontested line is disconnected. The default is 120 seconds.
Fast Idle Timeout	The default amount of idle time before a contested line is disconnected. The default is 20 seconds. Line contention occurs when a busy line is requested to send another packet to a different destination.

Dialer Physical Interface Dialog Box

Use the Dialer Physical Interface dialog box to add or edit the properties that associate physical BRI interfaces with dialer interfaces.

Note Use FlexConfigs to define other types of physical dialer interfaces, such as ATM and Ethernet. For more information, see Understanding FlexConfig Policies and Policy Objects, page 7-2.

Navigation Path

Go to the Dialer Policy Page, page 62-31, then click the **Add** or **Edit** button beneath the Dialer Physical Interfaces table.

Related Topics

- Dialer Profile Dialog Box, page 62-32
- Defining BRI Interface Properties, page 62-30
- Dialer Interfaces on Cisco IOS Routers, page 62-28
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 62-13 Dialer Physical Interface Dialog Box

Element	Description
ISDN BRI	The physical BRI interface associated with the dialer interface. Enter the name of an interface or interface role object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Pools	Associates dialer pools with a physical interface. Enter the names of one or more pools (as defined in the Dialer Profile Dialog Box, page 62-32), or click **Select** to display a selector. Use commas to separate multiple entries.

Table 62-13 Dialer Physical Interface Dialog Box (continued)

Element	Description
Switch Type	The ISDN switch type. Options for North America are: • basic-5ess—Lucent (AT&T) basic rate 5ESS switch • basic-dms100—Northern Telecom DMS-100 basic rate switch • basic-ni—National ISDN switches Options for Australia, Europe, and the UK are: • basic-1tr6—German 1TR6 ISDN switch • basic-net3—NET3 ISDN BRI for Norway NET3, Australia NET3, and New Zealand NET3 switch types; ETSI-compliant switch types for Euro-ISDN E-DSS1 signaling system • vn3—French VN3 and VN4 ISDN BRI switches Option for Japan is: • ntt—Japanese NTT ISDN switches Option for Voice/PBX system is: • basic-qsig—PINX (PBX) switches with QSIG signaling per Q.931 ()
SPID1	Applies only when you select Basic-DMS-100, Basic-NI, or Basic-5ess as the switch type. The service provider identifier (SPID) for the ISDN service to which the interface subscribes. Some service providers in North America assign SPIDs to ISDN devices when you first subscribe to an ISDN service. If you are using a service provider that requires SPIDs, your ISDN device cannot place or receive calls until it sends a valid assigned SPID to the service provider when accessing the switch to initialize the connection. Valid SPIDs can contain up to 20 characters, including spaces and special characters. **Note** We recommend that you do not enter a SPID for interfaces using the AT&T 5ESS switch type, even though they are supported.
SPID2	Applies only when you select DMS-100 or NI as the switch type. The service provider identifier (SPID) for a second ISDN service to which the interface subscribes. Valid SPIDs can contain up to 20 alphanumeric characters (no spaces are permitted).

ADSL on Cisco IOS Routers

Digital Subscriber Line (DSL) is a family of technologies that transports data over existing twisted-pair copper wire. DSL uses frequencies that are beyond the upper list used by POTS (plain old telephone service) to deliver broadband applications, such as multimedia and video, over the local loop (or *last mile*) that connects the telephone company's central office to customer sites.

Asymmetric Digital Subscriber Line (ADSL) is a form of DSL where the data flow downstream to customer sites is much greater than the data flow upstream to the central office (CO). This asymmetric setup is well-suited for applications where users typically download far more information than they send, such as web surfing, video-on-demand, and remote LAN access. With ADSL, the connection speed is related to the distance between the customer site and the digital subscriber line-access multiplexer (DSLAM) that aggregates the connections from multiple customer sites onto a high-speed line.

ADSL downstream rates range from 1.5 to 9 Mbps, whereas upstream bandwidth ranges from 16 to 640 kbps. ADSL transmissions work at distances up to 18,000 feet (5,488 meters) over a single copper twisted pair. Newer versions of ADSL technology, such as ADSL2 and ADSL2+, offer even higher data rates for short distances, as well as power management and realtime performance monitoring.

ATM is used in many ADSL implementations due to its small, fixed-length cell size, which makes it suitable for carrying time-critical traffic, such as voice and video, in conjunction with other traffic. You can use Security Manager to configure ATM over DSL on a Cisco IOS router. For more information about configuring ADSL policies in Security Manager, see Defining ADSL Settings, page 62-36.

To configure ADSL in Security Manager, you must do the following:

1. Configure an ATM interface or subinterface. See Defining Basic Router Interface Settings, page 62-4.
2. Configure ADSL settings on the ATM interface or subinterface. See Defining ADSL Settings, page 62-36.
3. Configure PVCs on the ATM interface or subinterface. See Defining ATM PVCs, page 62-51.

Note If you perform discovery on the device, Security Manager populates the Interfaces policy with the ATM interface and subinterface and the ADSL policy with the ADSL settings for that interface. Any discovered PVCs are added to the PVC policy.

Related Topics

- Supported ADSL Operating Modes, page 62-35

Supported ADSL Operating Modes

Table 62-14 on page 62-35 describes the operating modes that are supported on each ADSL interface card that can be configured with Security Manager.

Table 62-14 ADSL Cards and Supported DSL Operating Modes

ADSL Interface Card	Supported DSL Operating Modes
WIC-1ADSL	auto, ansi-dmt, itu-dmt, splitterless
WIC-1ADSL-I-DG	auto, etsi, itu-dmt
WIC-1ADSL-DG	auto, ansi-dmt, itu-dmt, splitterless
HWIC-1ADSL	auto, ansi-dmt, itu-dmt, adsl2, adsl2+
HWIC-1ADSLI	auto, etsi, itu-dmt, adsl2, adsl2+
HWIC-ADSL-B/ST	auto, ansi-dmt, itu-dmt, adsl2, adsl2+
HWIC-ADSLI-B/ST	auto, etsi, itu-dmt, adsl2, adsl2+

Table 62-15 on page 62-36 describes the operating modes that are supported on each ADSL device that can be configured with Security Manager.

Table 62-15 Fixed ADSL Devices and Supported DSL Operating Modes

Device	Supported DSL Operating Modes
857 Integrated Services Router	auto, ansi-dmt, itu-dmt, adsl2, adsl2+
876 Integrated Services Router	auto, etsi, itu-dmt, adsl2, adsl2+
877 Integrated Services Router	auto, ansi-dmt, itu-dmt, adsl2, adsl2+
1801 Integrated Services Router	auto, ansi-dmt, itu-dmt, adsl2, adsl2+
1802 Integrated Services Router	auto, etsi, itu-dmt, adsl2, adsl2+

Related Topics

- Defining ADSL Settings, page 62-36
- ADSL on Cisco IOS Routers, page 62-34

Defining ADSL Settings

When you configure an ADSL definition in Security Manager, you must select the ATM interface on which ADSL is being defined. In addition, we highly recommend that you specify the router type or the type of WIC (WAN interface card) installed in the router. The validity of DSL policy definitions is highly dependent on the hardware. By specifying the hardware used by this policy, you enable Security Manager to properly validate the values you define and avoid deployment failures.

You can optionally specify the following parameters:

- The DSL operating mode.
- Whether to enable dynamic VC bandwidth adjustments when using Inverse Multiplexing over ATM (IMA).
- Whether certain interface cards should use a particular set of carrier tones.

Modular Cisco IOS routers may contain multiple interface cards, each of which contains a single ATM interface. You may define only one ADSL definition per interface.

Before You Begin

- Make sure that the device contains an ADSL ATM interface. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Related Topics

- Supported ADSL Operating Modes, page 62-35
- ADSL on Cisco IOS Routers, page 62-34
- PVCs on Cisco IOS Routers, page 62-47

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > DSL > ADSL** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > DSL > ADSL** from the Policy Type selector. Select an existing policy or create a new one.

The ADSL page is displayed. See Table 62-16 on page 62-38 for a description of the fields on this page.

Step 2 Click the **Add** button beneath the table to display the ADSL Settings dialog box. See Table 62-17 on page 62-39 for a description of the fields in this dialog box.

Step 3 In the ATM Interface field, enter the name of the ATM interface or interface role on which you want to define ADSL settings, or click **Select** to select an interface role or create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Note The interface that you select must be physically present on the device; otherwise, deployment fails.

Step 4 (Optional) Select the interface card type installed on the router.

Note When discovering from a live device, the correct interface card type is already displayed. If you did not perform discovery on a live device, or if Security Manager cannot detect the type of interface card installed on the device, this field displays "Unknown".

Step 5 (Optional) When using IMA groups, select the **Allow bandwidth change on ATM PVCs** check box to enable dynamic adjustments to VC bandwidth in response to changes in group bandwidth. If this check box is left deselected, you must make these adjustments manually.

Step 6 (Optional) Specify the DSL operating mode for this ATM interface. See Table 62-14 on page 62-35 for a list of the operating modes supported for each card type.

Step 7 (Optional) Select the **Use low tone set** check box to have the interface card use carrier tones 29 through 48.

Step 8 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the ADSL table.

Note To edit an ADSL definition, select it from the table, then click **Edit**. To remove an ADSL definition, select it, then click **Delete**.

Step 9 Repeat Step 2 through Step 8 to define ADSL settings on additional ATM interfaces. Only one ADSL definition may be defined on an interface.

ADSL Policy Page

Use the ADSL page to create, edit, and delete ADSL definitions on the ATM interfaces of the router. For more information, see Defining ADSL Settings, page 62-36.

Navigation Path

- (Device view) Select **Interfaces > Settings > DSL > ADSL** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > DSL > ADSL** from the Policy Type selector. Right-click **ADSL** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- PVC Policy Page, page 62-55
- SHDSL Policy Page, page 62-42
- ADSL on Cisco IOS Routers, page 62-34
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-16 ADSL Page

Element	Description
ATM Interface	The ATM interface on which ADSL settings are defined.
Interface Card	The type of device or ADSL interface card on which the ATM interface resides.
Bandwidth Change	Indicates whether the router makes dynamic adjustments to VC bandwidth as overall bandwidth changes. (This is relevant only when IMA groups are configured on the ATM interface.)
DSL Operating Mode	The DSL operating mode for this interface.
Tone Low	Indicates whether the interface is using the low tone set (carrier tones 29 through 48).
Add button	Opens the ADSL Settings Dialog Box, page 62-38. From here you can define the ADSL settings for a selected ATM interface.
Edit button	Opens the ADSL Settings Dialog Box, page 62-38. From here you can edit the selected ADSL definition.
Delete button	Deletes the selected ADSL definition from the table.

ADSL Settings Dialog Box

Use the ADSL Settings dialog box to configure ADSL settings on a selected ATM interface.

Note When you configure ADSL settings, we highly recommend that you select the type of device or interface card on which the ATM interface is defined. ADSL settings are highly dependent on the hardware. Defining the hardware type in Security Manager enables proper validation of your configuration for a successful deployment to your devices.

Navigation Path

Go to the ADSL Policy Page, page 62-37, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining ADSL Settings, page 62-36
- PVC Policy Page, page 62-55

Field Reference

Table 62-17 ADSL Settings Dialog Box

Element	Description
ATM Interface	The ATM interface on which ADSL settings are defined. Enter the name of an interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. **Note** We recommend that you do not define an interface role that includes ATM interfaces from different interface cards. The different settings supported by each card type may cause deployment to fail. **Note** You can create only one ADSL definition per interface.
Interface Card	The device type or the type of interface card installed on the router: • [blank]—The interface card type is not defined. • WIC-1ADSL—A 1-port ADSL WAN interface card that provides ADSL over POTS (ordinary telephone lines). • WIC-1ADSL-I-DG—A 1-port ADSL WAN interface card that provides ADSL over ISDN with Dying Gasp support. (With Dying Gasp, the router warns the DSLAM of imminent line drops when the router is about to lose power.) • WIC-1ADSL-DG—A 1-port ADSL WAN interface card that provides ADSL over POTS with Dying Gasp support. • HWIC-1ADSL—A 1-port high-speed ADSL WAN interface card that provides ADSL over POTS. • HWIC-1ADSLI—A 1-port high-speed ADSL WAN interface card that provides ADSL over ISDN. • HWIC-ADSL-B/ST—A 2-port high-speed ADSL WAN interface card that provides ADSL over POTS with an ISDN BRI port for backup. • HWIC-ADSLI-B/ST—A 2-port high-speed ADSL WAN interface card that provides ADSL over ISDN with an ISDN BRI port for backup.

Table 62-17 ADSL Settings Dialog Box (continued)

Element	Description
Interface Card (continued)	• 857 ADSL—Cisco 857 Integrated Service Router with an ADSL interface. • 876 ADSL—Cisco 876 Integrated Services Router with an ADSL interface. • 877 ADSL—Cisco 877 Integrated Services Router with an ADSL interface. • 1801 ADSLoPOTS—Cisco 1801 Integrated Services Router that provides ADSL over POTS. • 1802 ADSLoISDN—Cisco 1802 Integrated Services Router that provides ADSL over ISDN. **Note** When discovering from a live device, the correct interface card type will already be displayed. If you did not perform discovery on a live device, or if Security Manager cannot detect the type of interface card installed on the device, this field displays "Unknown".
Allow bandwidth change on ATM PVCs	When selected, the router makes dynamic adjustments to VC bandwidth in response to changes in the overall bandwidth of the Inverse Multiplexing over ATM (IMA) group defined on the ATM interface. When deselected, PVC bandwidth must be adjusted manually (using the CLI) whenever an individual physical link in the IMA group goes up or down.
DSL Operating Mode	The operating mode configured for this ADSL line: • auto—Performs automatic negotiation with the DSLAM located at the central office (CO). This is the default. • ansi-dmt—The line trains in ANSI T1.413 Issue 2 mode. • itu-dmt—The line trains in G.992.1 mode. • splitterless—The line trains in G.992.2 (G.Lite) mode. • etsi—The line trains in ETSI (European Telecommunications Standards Institute) mode. • adsl2—The line trains in G.992.3 (adsl2)mode. • adsl2+—The line trains in G.992.5 (adsl2+) mode. **Note** See Table 62-14 on page 62-35 for a description of the operating modes that are supported by each card type.
Use low tone set	When selected, the interface card uses carrier tones 29 through 48. When deselected, the interface card uses carrier tones 33 through 56. **Note** Leave this option deselected when the interface card is operating in accordance with Deutsche Telekom specification U-R2.

SHDSL on Cisco IOS Routers

Digital Subscriber Line (DSL) is a family of technologies that transports data over existing twisted-pair copper wire. DSL uses frequencies that are beyond the upper list used by POTS (plain old telephone service) to deliver broadband applications, such as multimedia and video, over the local loop (or *last mile*) that connects the telephone company's central office to customer sites.

Based on the International Telecommunications Union (ITU) G.991.2 global industry standard, symmetric high-speed digital subscriber line (SHDSL) delivers symmetrical data rates from 192 up to 2.3 Mbps on a single wire pair. It transports many types of signals, such as T1, E1, ISDN, ATM, and IP. In addition, the G.SHDSL signal has a greater distance reach from the central office than ADSL and proprietary SDSL connections.

To configure SHDSL in Security Manager, do the following:

1. Configure the SHDSL controller. See Defining SHDSL Controllers, page 62-41.
2. Deploy the SHDSL policy. If ATM mode is activated, the router creates an ATM interface that corresponds to the controller upon deployment. See Working with Deployment and the Configuration Archive, page 8-25.
3. Rediscover the device to add the new ATM interface to Security Manager. See Discovering Policies on Devices Already in Security Manager, page 5-15.
4. (Optional) Create one or more subinterfaces on the ATM interface. See Defining Basic Router Interface Settings, page 62-4.
5. Configure PVCs on the ATM interface or subinterface. See Defining ATM PVCs, page 62-51.

Note If you perform discovery on the device, Security Manager populates the SHDSL policy with the definition of the controller and the Interfaces policy with the ATM interface and subinterface. Any discovered PVCs are added to the PVC policy.

Related Topics

- PVCs on Cisco IOS Routers, page 62-47

Defining SHDSL Controllers

When you configure an SHDSL controller in Security Manager, you must enter the name of the controller that is installed in the Cisco IOS router. The following settings are then applied automatically:

- ATM mode is enabled.
- The line termination is set to CPE (customer premises equipment).
- The line mode is set to Auto.

You can optionally change the line termination to CO and specify the DSL mode and line mode. In addition, you can define signal-to-noise ratio margins to improve line stability.

A Cisco IOS router may contain multiple SHDSL controllers. You may define only one SHDSL definition per controller.

Note When you deploy an SHDSL policy with ATM mode enabled, an ATM interface is created automatically on the router. Perform rediscovery to add the interface into Security Manager. You can then define PVCs on the ATM interface as required. See Defining ATM PVCs, page 62-51.

Before You Begin

- Make sure that an SHDSL controller in installed on the device.

Related Topics

- SHDSL on Cisco IOS Routers, page 62-41
- PVCs on Cisco IOS Routers, page 62-47

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > DSL > SHDSL** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > DSL > SHDSL** from the Policy Type selector. Select an existing policy or create a new one.

The SHDSL page is displayed. See Table 62-18 on page 62-43 for a description of the fields on this page.

Step 2 Click the **Add** button beneath the table to display the SHDSL dialog box.

Step 3 Enter the name of the controller, or click **Select** to display the utility for generating the name. See Controller Auto Name Generator Dialog Box, page 62-46.

Note The controller that you select must be physically present on the device; otherwise, deployment fails.

Step 4 Define the SHDSL controller as required. For more information, see Table 62-19 on page 62-44.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the SHDSL table.

Note To edit an SHDSL controller, select it from the table, then click **Edit**. To remove an SHDSL controller, select it, then click **Delete**.

Step 6 Repeat Step 2 through Step 5 to define additional SHDSL controllers. Only one definition may be defined per controller.

SHDSL Policy Page

Use the SHDSL page to create, edit, and delete DSL controller definitions on the router. For more information, see Defining SHDSL Controllers, page 62-41.

Navigation Path

- (Device view) Select **Interfaces > Settings > DSL > SHDSL** from the Policy selector.

- (Policy view) Select **Router Interfaces > Settings > DSL > SHDSL** from the Policy Type selector. Right-click **SHDSL** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- PVC Policy Page, page 62-55
- ADSL Policy Page, page 62-37
- SHDSL on Cisco IOS Routers, page 62-41
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-18 SHDSL Page

Element	Description
Name	The name of the DSL controller.
Description	An optional description of the controller.
Shutdown	Indicates whether the DSL controller is in shutdown mode.
Configure ATM Mode	Indicates whether the DSL controller has been set into ATM mode.
Line Termination	The line termination set for the router (CPE or CO).
DSL Mode	The operating mode defined for the DSL controller.
Line Mode	The line mode defined for the DSL controller.
Line Rate	The line rate (in kbps) defined for the DSL controller. **Note** A value is displayed in this column only if the line mode is not set to Auto.
SNR Margin Current	The current signal-to-noise ratio on the controller.
SNR Margin Snext	The self near-end crosstalk (Snext) signal-to-noise ratio on the controller.
Add button	Opens the SHDSL Controller Dialog Box, page 62-43. From here you can define the settings for a DSL controller.
Edit button	Opens the SHDSL Controller Dialog Box, page 62-43. From here you can edit the selected DSL controller definition.
Delete button	Deletes the selected DSL controller definition from the table.

SHDSL Controller Dialog Box

Use the SHDSL Controller dialog box to configure SHDSL controllers.

Navigation Path

Go to the SHDSL Policy Page, page 62-42, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining SHDSL Controllers, page 62-41
- PVC Policy Page, page 62-55

- Discovering Policies on Devices Already in Security Manager, page 5-15

Field Reference

Table 62-19 SHDSL Dialog Box

Element	Description
Name	The name of the controller. Enter a name manually, or click **Select** to display a dialog box for generating a name. See Controller Auto Name Generator Dialog Box, page 62-46.
Description	Additional information about the controller (up to 80 characters).
Shutdown	When selected, the DSL controller is in shutdown state. However, its definition is not deleted. When deselected, the DSL controller is enabled. This is the default.
Configure ATM mode	When selected, sets the controller into ATM mode and creates an ATM interface with the same ID as the controller. This is the default. You must enable ATM mode and then perform rediscovery to configure ATM or PVCs on the device. When deselected, ATM mode is disabled. No ATM interface is created on deployment. **Note** You cannot remove ATM mode from a controller after it has been saved in Security Manager.
Line Termination	The line termination that is set for the router: • CPE—Customer premises equipment. This is the default. • CO—Central office.
DSL Mode	The DSL operating mode, including regional operating parameters, used by the controller: • [blank]—The operating mode is not defined. (When deployed, the Annex A standard for North America is used.) • A—Supports Annex A of the G.991.2 standard for North America. • A-B—Supports Annex A or Annex B. Available only when the Line Term is set to CPE. The appropriate mode is selected when the line trains. • A-B-ANFP—Supports Annex A or Annex B-ANFP. Available only when the Line Term is set to CPE. The appropriate mode is selected when the line trains. • B—Supports Annex B of the G.991.2 standard for Europe. • B-ANFP—Supports Annex B-ANFP (Access Network Frequency Plan). **Note** The available DSL modes are dependent on the selected line termination.
Line Mode settings	

Table 62-19 SHDSL Dialog Box (continued)

Element	Description
Line Mode	The line mode used by the controller: • auto—The controller operates in the same mode as the other line termination (2-wire line 0, 2-wire line 1, or 4-wire enhanced). This is the default for CPE line termination. • 2-wire—The controller operates in two-wire mode. This is the default for CO line termination. • 4-wire—The controller operates in four-wire mode. **Note** You can select Auto only when you configure the controller as the CPE.
Line	Applies only when the Line Mode is defined as 2-wire. The pair of wires to use: • line-zero—RJ-11 pin 1 and pin 2. This is the default for CO line termination. • line-one—RJ-11 pin 3 and pin 4.
Exchange Handshake	Applies only when the Line Mode is defined as 4-wire. The type of handshake mode to use: • [blank]—The handshake mode is not specified. (When deployed, the enhanced option is used.) This is the default. • enhanced—Exchanges handshake status on both wire pairs. • standard—Exchanges handshake status on the main wire pair only.
Line Rate	Does not apply when the Line Mode is defined as Auto. The DSL line rate (in kbps) available for the SHDSL port: • auto—The controller selects the line rate. This is available only in 2-wire mode. • Supported line rates: – For 2-wire mode: 192, 256, 320, 384, 448, 512, 576, 640, 704, 768, 832, 896, 960, 1024, 1088, 1152, 1216, 1280, 1344, 1408, 1472, 1536, 1600, 1664, 1728, 1792, 1856, 1920, 1984, 2048, 2112, 2176, 2240, and 2304. – For 4-wire mode: 384, 512, 640, 768, 896, 1024, 1152, 1280, 1408, 1536, 1664, 1792, 1920, 2048, 2176, 2304, 2432, 2560, 2688, 2816, 2944, 3072, 3200, 3328, 3456, 3584, 3712, 3840, 3968, 4096, 4224, 4352, 4480, and 4608. **Note** Third-party equipment may use a line rate that includes an additional SHDSL overhead of 8 kbps for 2-wire mode or 16 kbps for 4-wire mode.
SNR Margin settings	

Table 62-19 SHDSL Dialog Box (continued)

Element	Description
Current	The current signal-to-noise (SNR) ratio on the controller, in decibels (dB). Valid values range from -10 to 10 dB. This option can create a more stable line by making the line train more than current noise margin plus SNR ratio threshold during training time. If any external noise is applied that is less than the set SNR margin, the line will be stable. **Note** Select **disable** to disable the current SNR.
Snext	The Self Near-End Crosstalk (SNEXT) signal-to-noise ratio on the controller, in decibels. Valid values range from -10 to 10 dB. This option can create a more stable line by making the line train more than SNEXT threshold during training time. If any external noise is applied that is less than the set SNEXT margin, the line will be stable. **Note** Select **disable** to disable the SNEXT SNR.

Controller Auto Name Generator Dialog Box

Use the Controller Auto Name Generator dialog box to have Security Manager generate a name for the DSL controller based on its location in the router.

Navigation Path

Go to the SHDSL Controller Dialog Box, page 62-43, then click **Select** in the Name field**.**

Related Topics

- Defining SHDSL Controllers, page 62-41
- SHDSL Policy Page, page 62-42
- PVC Policy Page, page 62-55

Field Reference

Table 62-20 Controller Auto Name Generator Dialog Box

Element	Description
Type	The type of interface. This field displays the value DSL and is read-only.
Card	The card related to the controller.
Slot	The slot related to the controller.
Port	The port related to the controller. **Note** The information you enter in these fields forms the remainder of the generated name, as displayed in the Result field.

Table 62-20 Controller Auto Name Generator Dialog Box (continued)

Element	Description
Result	The name generated by Security Manager from the information you entered for the controller location. The name displayed in this field is read-only. **Tip** After closing this dialog box, you can edit the generated name in the SHDSL dialog box, if required.

PVCs on Cisco IOS Routers

Asynchronous Transfer Mode (ATM) is an International Telecommunication Union (ITU-T) standard designed for the high-speed transfer of voice, video, and data through public and private networks using cell relay technology. A cell switching and multiplexing technology, ATM combines the benefits of circuit switching (constant transmission delay, guaranteed capacity) with those of packet switching (flexibility, efficiency for intermittent traffic). An ATM network is made up of one or more ATM switches and ATM endpoints, such as a Cisco IOS router.

There are three general types of ATM services, permanent virtual connections (PVCs), switched virtual connections (SVCs), and connectionless service. PVCs allow direct and permanent connections between sites to provide a service that is similar to a leased line. Advantages of PVCs are the guaranteed availability of a connection and that no call setup procedures are required between switches. Each piece of equipment between the source and destination must be manually provisioned for the PVC.

For more information about ATM PVCs, see:

- Understanding Virtual Paths and Virtual Channels, page 62-47
- Understanding ATM Service Classes, page 62-48
- Understanding ATM Management Protocols, page 62-49

For more information about defining PVCs in Security Manager, see:

- Defining ATM PVCs, page 62-51
- Defining OAM Management on ATM PVCs, page 62-54

Related Topics

- ADSL on Cisco IOS Routers, page 62-34
- SHDSL on Cisco IOS Routers, page 62-41

Understanding Virtual Paths and Virtual Channels

ATM networks are fundamentally connection oriented. This means that a virtual connection needs to be established across the ATM network before any data transfer. Two types of ATM connections exist:

- Virtual path connections (VPCs), identified by a virtual path identifier (VPI).
- Virtual channel connections (VCCs), identified by the combination of a VPI and a VCI (virtual channel identifier). PVCs are a type of VCC where a permanent connection is defined between two sites.

As shown in Figure 62-2, a virtual path is a bundle of virtual channels, all of which are switched transparently across the ATM network on the basis of the common VPI. A VPC can be thought of as a bundle of VCCs with the same VPI value.

Figure 62-2 ATM Virtual Path and Virtual Channel Connections

Every cell header contains a VPI field and a VCI field, which explicitly associate a cell with a given virtual channel on a physical link. It is important to remember the following attributes of VPIs and VCIs:

- VPIs and VCIs are not addresses, such as MAC addresses used in LAN switching.
- VPIs and VCIs are explicitly assigned at each segment of a connection and, as such, have only local significance across a particular link. They are remapped, as appropriate, at each switching point.

Using the VPI/VCI identifier, the ATM layer can multiplex (interleave), demultiplex, and switch cells from multiple connections. Certain VPI/VCI identifiers are reserved for particular uses, such as the Integrated Local Management Interface (ILMI).

Related Topics

- Understanding ATM Service Classes, page 62-48
- Understanding ATM Management Protocols, page 62-49
- Defining ATM PVCs, page 62-51
- PVCs on Cisco IOS Routers, page 62-47

Understanding ATM Service Classes

Version 4.0 of the Traffic Management Specification published by the ATM Forum defines five service classes that describe the user traffic transmitted on a network and the quality of service that a network needs to provide for that traffic. Security Manager supports the following ATM service classes:

- *Available Bit Rate (ABR)* This is a service class where ATM switches make no guarantee of cell delivery, but do guarantee a minimum bit rate and that cell loss is kept as low as possible with the use of a feedback mechanism. The ABR service category is designed for VCs that carry file transfers and other bursty, non-real-time traffic that requires a minimum amount of bandwidth. This bandwidth is specified via a minimum cell rate that must be available while the VC is configured and active. For more details, see *Understanding the Available Bit Rate (ABR) Service Category for ATM VCs* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a00800fbc76.shtml.
- *Constant Bit Rate (CBR)* This is a service class where cells are transmitted in a continuous bitstream to meet voice and video QoS needs. The CBR service class is designed for ATM virtual circuits (VCs) that need a static amount of bandwidth that is continuously available for the duration of the active connection. An ATM VC configured as CBR can send cells at peak cell rate (PCR) at any time and for any duration. It also can send cells at a rate less than the PCR or even emit no cells. The

configuration on CBR may vary with different platforms. For more details, see *Understanding the CBR Service Category for ATM VCs* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a0080094e6a.shtml.

- *Unspecified Bit Rate (UBR)* This is a service class where the network management makes no Quality of Service (QoS) commitment. It models the best-effort service that the Internet normally provides and is suitable for applications tolerant to delay that do not require real-time responses. Examples include email, fax transmission, file transfers, Telnet, LAN and remote office interconnections. For more details, see *Understanding the UBR Service Category for ATM Virtual Circuits* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a00800a4837.shtml.
- *Unspecified Bit Rate (UBR+)* Cisco provides a variant of the UBR service class called UBR+. The main advantage of the UBR+ service class is that it allows an ATM end-system to signal a minimum cell rate to an ATM switch in a connection request, and the ATM network attempts to maintain this minimum as an end-to-end guarantee. For more details, see *Understanding the UBR+ Service Category for ATM VCs* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a0080094b40.shtml.
- *Variable Bit Rate - Non-Real Time (VBR-nrt)* This service class is used to transmit non-real-time applications that are bursty in nature. The traffic characteristics are defined in terms of the Peak Cell Rate (PCR), Sustained Cell Rate (SCR), and Minimum Burst Size (MBS). For more details, see *Understanding the VBR-nrt Service Category and Traffic Shaping for ATM VCs* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a0080102a42.shtml.
- *Variable Bit Rate - Real Time (VBR-rt)* This service class is used to transmit real-time data that is sensitive to time delays, like compressed voice over IP and video conferencing. As with VBR-nrt, VBR-rt traffic is defined in terms of a PCR, SCR, and MBS. For more details, see *Understanding the Variable Bit Rate Real Time (VBR-rt) Service Category for ATM VCs* at: http://www.cisco.com/en/US/tech/tk39/tk51/technologies_tech_note09186a0080094cd0.shtml.

You can use these service classes to define ATM quality of service (QoS) guarantees, such as traffic shaping. Traffic shaping is the use of queues to constrain data bursts, limit peak data rate, and smooth jitter so that traffic fits within the envelope defined by the traffic contract. ATM devices use traffic shaping to adhere to the terms of the traffic contract.

Related Topics

- Understanding Virtual Paths and Virtual Channels, page 62-47
- Understanding ATM Management Protocols, page 62-49
- Defining ATM PVCs, page 62-51
- PVCs on Cisco IOS Routers, page 62-47

Understanding ATM Management Protocols

ATM uses two different types of signaling for tracking the status of PVCs:

- Integrated Local Management Interface (ILMI). For more information, see Understanding ILMI, page 62-50.
- Flow 4 (F4) and Flow 5 (F5) Operation, Administration, and Maintenance (OAM) cells. For more information, see Understanding OAM, page 62-51.

Security Manager can be used to enable and disable ILMI on specific PVCs and to configure F5 OAM functionality.

Related Topics

- Understanding Virtual Paths and Virtual Channels, page 62-47
- Understanding ATM Service Classes, page 62-48
- Defining ATM PVCs, page 62-51
- Defining OAM Management on ATM PVCs, page 62-54
- PVCs on Cisco IOS Routers, page 62-47

Understanding ILMI

The Integrated Local Management Interface (ILMI) is a protocol defined by the ATM Forum for setting and capturing physical layer, ATM layer, virtual path, and virtual circuit parameters on ATM interfaces. ILMI facilitates network-wide autoconfiguration by enabling devices to determine the status of components at the other end of a physical link and to negotiate a common set of operational parameters to ensure interoperability. The ATM routing protocols, PNNI (Private Network to Network Interface) and IISP (Interim-Interswitch Signaling Protocol), use this information to discover and bring up a network of interconnected ATM switch routers.

When two ATM interfaces run the ILMI protocol, they exchange ILMI packets across the physical connection. These packets consist of SNMP messages as large as 484 octets. ATM interfaces encapsulate these messages in an ATM adaptation layer 5 (AAL5) trailer, segment the packet into cells, and schedule the cells for transmission. ATM interfaces use the SNMP object IDs in network functions such as permanent virtual circuit (PVC) autodiscovery, which is particularly useful in digital subscriber line (DSL) applications.

ILMI organizes managed objects into multiple information bases (MIBs), including one for link management. This MIB contains the following object groups for all ATM interfaces:

- Physical layer—ILMI 4.0 discontinues or "deprecates" earlier physical-layer ILMI values and specifies the use of the standard Interface MIB (RFC 1213).
- ATM layer—Indicates the number of available bits for VPI and VCI values in the ATM cell header, maximum number of virtual path connections (VPCs) and virtual channel connections (VCCs) allowed, number of configured PVCs, and so on.
- Virtual path connection—Indicates the up or down status of a VPC and its Quality of Service (QoS) parameters.
- Virtual channel connection—Indicates the up or down status of the VCC and its QoS parameters.

Administrators may enable or disable ILMI at will, but we highly recommend you enable it. Without ILMI, you must manually configure many of the parameters otherwise managed by ILMI for the ATM devices to operate correctly. ILMI operates over a reserved PVC of VPI=X, VCI=16.

Related Topics

- Understanding ATM Management Protocols, page 62-49
- PVCs on Cisco IOS Routers, page 62-47

Understanding OAM

The Operation, Administration, and Maintenance (OAM) feature provides fault management and performance management for ATM and is based on the standard defined in ITU recommendation I.610. OAM detects network connectivity failures on a PVC and reacts by bringing down the PVC. Without OAM, a PVC would remain up after network connectivity is lost. In such a situation, routing table entries would continue to point to the PVC, resulting in lost packets.

Security Manager enables the use of F5 OAM, which operates at the virtual circuit (VC) level. To detect a failure along the PVC path on an end-device, such as a Cisco IOS router, OAM uses the following cells:

- Loopback cells—At regular intervals, routers configured for OAM send loopback cells which must be looped in the network. This looping point can be the machine at the end of the PVC (end-to-end loopback cells) or a device on the path (segment loopback cells). A failure occurs when the loopback cell fails to return to its point of origin.
- Continuity Check (CC) cells—CC cells are sent regularly by routers configured for OAM to check the integrity of the link. CC cells can be sent either end-to-end or confined to a particular segment of the PVC. Activation and deactivation cells are used to initiate and suspend continuity checking. Any connectivity failures are reported in special SNMP notifications.
- Alarm Indication Signal (AIS) cells—In the event of a failure at the physical layer, AIS cells are sent to downstream devices to report a virtual connection failure at the ATM layer. The PVC moves to the down state after a defined number of AIS cells are received and does not come up again until a defined interval passes without additional AIS cells.
- Remote Detection Indication (RDI) cells—When AIS cells are sent to warn downstream devices of a connectivity failure, RDI cells are sent upstream in response as a control and feedback mechanism for the network.

AIS/RDI cells are sent using the same VPI/VCI as the user cells on the affected PVC until the failure is resolved.

Related Topics

- Understanding ATM Management Protocols, page 62-49
- PVCs on Cisco IOS Routers, page 62-47
- Defining OAM Management on ATM PVCs, page 62-54

Defining ATM PVCs

You define an ATM permanent virtual circuit (PVC) by selecting an ATM interface and then defining the following settings:

- The PVC ID.
- The type of encapsulation to use.
- Whether ILMI management is enabled on this PVC.
- Whether Inverse ARP (InARP) is used to learn the IP addresses of the destination devices.
- Options related to PPP over Ethernet (PPPoE) and PPP over ATM (PPPoA).
- Quality-of-service settings, such as traffic shaping.
- Static IP address mappings in place of InARP.

For information about defining F5 Operation, Administration, and Maintenance (OAM) management, such as loopbacks and continuity checks, on PVCs, see Defining OAM Management on ATM PVCs, page 62-54.

Before You Begin

- When configuring ATM over DSL, make sure that you have configured either an ADSL policy (seeADSL on Cisco IOS Routers, page 62-34) or an SHDSL policy (SHDSL on Cisco IOS Routers, page 62-41).
- Make sure that the device contains an ATM interface and subinterface. (PVCs are typically configured on ATM subinterfaces.) See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Note When configuring ATM for SHDSL, the ATM interface is created when you define the SHDSL controller and enable ATM mode. You must then rediscover the device to add the ATM interface to Security Manager. See Defining SHDSL Controllers, page 62-41.

Related Topics

- Defining OAM Management on ATM PVCs, page 62-54
- Understanding Policing and Shaping Parameters, page 66-6
- PVCs on Cisco IOS Routers, page 62-47

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > PVC** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > PVC** from the Policy Type selector. Select an existing policy or create a new one.

The PVC page is displayed. See Table 62-21 on page 62-56 for a description of the fields on this page.

Step 2 Click the **Add** button beneath the table to display the PVC dialog box. See Table 62-22 on page 62-57 for a description of the fields in this dialog box.

Step 3 In the Interface field, enter the name of the ATM interface, ATM subinterface, or interface role on which you want to define the PVC, or click **Select** to select an interface role or to create a new one.

Step 4 Select the type of device or DSL WAN interface card that contains the ATM interface.

Note We highly recommend that you define this setting to ensure the proper validation of your PVC policy, as the settings in this policy are highly hardware-dependent.

Step 5 On the Settings tab of the PVC dialog box, define the basic settings of the PVC:

a. Enter the VPI/VCI identifier and an optional text handle. If you are defining the management PVC, select the **Management PVC (ILMI)** check box.

Note An error occurs if two users attempt to define PVCs with the same identifiers at the same time.

b. Select the type of ATM encapsulation to use. If you select aal5autoppp or aal5ciscoppp, you must define the virtual template to use for PPPoA, or click Select to display a selector. If you select aal5mux as the encapsulation type, you must select the protocol that is carried by the PVC.

Note Do not select an encapsulation type when defining the management PVC.

Note If you modify the virtual template settings on an existing PVC, you must enter the **shutdown** command followed by the **no shutdown** command on the ATM subinterface to restart the interface. This causes the newly configured parameters to take effect.

c. Select the Enable ILMI check box to enable the ILMI to manage this PVC. For more information, see Understanding ILMI, page 62-50.

Note You cannot configure the management PVC on a subinterface.

d. Select the Inverse ARP check box to enable the PVC to dynamically learn the Layer 3 addresses that are required to forward traffic to those devices.

Note Alternatively, you can create static address mappings, as described in Step 7.

e. In the PPPoE Max Sessions field, define the maximum number of PPPoE sessions allowed on the PVC.

f. In the VPN Service Name field, define the static domain name to use for PPPoA sessions on the PVC.

See Table 62-23 on page 62-58for a description of the fields on the Settings tab.

Step 6 (Optional) On the QoS tab of the PVC dialog box, define the type of ATM traffic shaping to perform on the traffic carried by this PVC. Traffic shaping regulates the flow of traffic carried by the PVC by queuing traffic that exceeds the defined bit rates. See Table 62-24 on page 62-62 for a description of the fields on the QoS tab.

Step 7 (Optional) On the Protocol tab of the PVC dialog box, create static mappings for the IP addresses at the other end of the PVC:

a. Click **Add** to display the Define Mapping dialog box. See Table 62-26 on page 62-66 for a description of the fields in this dialog box.

b. Select IP Address, then enter the address or network/host object that you want to map, or click **Select** to select a network/host object from a list or to create a new one.

c. Click **OK**. The static mapping is displayed on the Protocol tab.

d. Repeat a. through c. to define additional static mappings.

Note You can also use the Protocol tab to change the type of InARP to use, broadcast or non-broadcast.

Step 8 Click **Advanced** to configure OAM management on the PVC. See Defining OAM Management on ATM PVCs, page 62-54.

Step 9 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the PVC table.

Note To edit a PVC, select it from the table, then click **Edit**. To remove a PVC, select it, then click **Delete**.

Step 10 Repeat Step 2 through Step 9 to define additional PVCs.

Defining OAM Management on ATM PVCs

Security Manager enables you to configure the following F5 (VC level) Operation, Administration, and Maintenance (OAM) cells for detecting PVC failures in a Cisco IOS router:

- Loopback cells
- Continuity Check (CC) cells
- Alarm Indication Signal (AIS) cells
- Remote Detection Indication (RDI) cells

You can enable and disable each of these cell types and define settings that determine how each cell type affects the PVC when a failure is detected.

Before You Begin

- Select the ATM interface on which the PVC is defined.
- Define the general settings and the QoS settings of the PVC. See Defining ATM PVCs, page 62-51.

Related Topics

- Defining ATM PVCs, page 62-51
- PVCs on Cisco IOS Routers, page 62-47

Step 1 In the PVC dialog box, click **Advanced** to display the PVC Advanced Settings dialog box. See Table 62-27 on page 62-66 for a description of the fields in this dialog box.

Step 2 Enable OAM loopback cells on the selected PVC:

a. Click the **OAM-PVC** tab. See Table 62-29 on page 62-69 for a description of the fields on this tab.

b. Select the **Enable OAM Management** check box.

c. Define the frequency of loopback cell transmissions.

Step 3 (Optional) Enable segment CC cells on the PVC:

a. Under Segment Continuity Check, select **Configure Continuity Check**.

b. Choose whether the router should act as the sink, source, or both. This determines the direction in which CC cells are sent.

c. Choose whether the PVC should remain up after segment or end-to-end failures are detected.

Note Select **Deny Activation Requests** to have the router reject CC activation requests received from peers.

Step 4 (Optional) Enable end-to-end CC cells on the PVC, using the procedure described in Step 3 for segment CC cells.

Step 5 (Optional) Configure additional loopback cell parameters:

a. Click the **OAM** tab.

b. Select the **Enable OAM Retry** check box, then define the down count, up count, and retry frequency. See Table 62-28 on page 62-67 for a description of the available options.

Step 6 (Optional) Configure additional CC cell parameters:

a. Select the **Enable** check box for segment CC cells, then define the activation count, deactivation count, and retry frequency. These fields determine how many activation and deactivation requests are sent to peers and how often the router waits between each attempt. See Table 62-28 on page 62-67 for a description of the available options.

b. Repeat a. for end-to-end CC cells.

Step 7 (Optional) Configure AIS/RDI cells on the PVC:

a. On the OAM tab, select the **Enable AIS-RDI Detection** check box.

b. Define how many AIS/RDI cells are required to move the PVC to the down state.

c. Define how many seconds must elapse without receiving AIS/RDI cells before moving the PVC to the up state.

Step 8 Click **OK** to close the dialog box and return to PVC dialog box.

PVC Policy Page

Use the PVC page to create, edit, and delete permanent virtual connections (PVCs) on the router. PVCs allow direct and permanent connections between sites to provide a service that is similar to a leased line. These PVCs can be used in ADSL, SHDSL, or pure ATM environments. For more information, see Defining ATM PVCs, page 62-51.

Navigation Path

- (Device view) Select **Interfaces > Settings > PVC** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > PVC** from the Policy Type selector. Right-click **PVC** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- ADSL Policy Page, page 62-37
- SHDSL Policy Page, page 62-42
- PVCs on Cisco IOS Routers, page 62-47
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-21 PVC Page

Element	Description
ATM Interface	The ATM interface on which the PVC is defined.
Interface Card	The type of device or WAN interface card on which the ATM interface resides.
PVC ID	The Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) of the PVC.
Settings	Additional settings configured for the PVC, including encapsulation, the number of PPPoE sessions, and the VPN service name.
QoS	Quality-of-service settings defined for the PVC, such as traffic shaping.
Protocol	The IP protocol mappings (static maps or Inverse ARP) configured for the PVC.
OAM	The F5 Operation, Administration, and Maintenance (OAM) loopback, continuity check, and AIS/RDI definitions configured for the PVC.
OAM-PVC	The OAM management cells that are configured for the PVC.
Add button	Opens the PVC Dialog Box, page 62-56. From here you can define a PVC.
Edit button	Opens the PVC Dialog Box, page 62-56. From here you can edit the selected PVC.
Delete button	Deletes the selected PVC from the table.

PVC Dialog Box

Use the PVC dialog box to configure ATM permanent virtual circuits (PVCs).

You can configure the following types of interface cards:

- Unknown—The interface card type is not defined.
- WIC-1ADSL—A 1-port ADSL WAN interface card that provides ADSL over POTS (ordinary telephone lines).
- WIC-1ADSL-I-DG—A 1-port ADSL WAN interface card that provides ADSL over ISDN with Dying Gasp support. (With Dying Gasp, the router warns the DSLAM of imminent line drops when the router is about to lose power.)
- WIC-1ADSL-DG—A 1-port ADSL WAN interface card that provides ADSL over POTS with Dying Gasp support.
- HWIC-1ADSL—A 1-port high-speed ADSL WAN interface card that provides ADSL over POTS.
- HWIC-1ADSLI—A 1-port high-speed ADSL WAN interface card that provides ADSL over ISDN.
- HWIC-ADSL-B/ST—A 2-port high-speed ADSL WAN interface card that provides ADSL over POTS with an ISDN BRI port for backup.
- HWIC-ADSLI-B/ST—A 2-port high-speed ADSL WAN interface card that provides ADSL over ISDN with an ISDN BRI port for backup.
- WIC-1-SHDSL-V2—A 1-port multiline G.SHDSL WAN interface card with support for 2-wire mode and enhanced 4-wire mode.

- WIC-1-SHDSL-V3—A 1-port multiline G.SHDSL WAN interface card with support for 2-wire mode and 4-wire mode (standard & enhanced).
- NM-1A-T3—A 1-port ATM network module with a T3 link.
- NM-1A-OC3-POM—A 1-port ATM network module with an optical carrier level 3 (OC-3) link and three operating modes (multimode, single-mode intermediate reach (SMIR), and single-mode long-reach (SMLR)).
- NM-1A-E3—A 1-port ATM network module with an E3 link.
- 857 ADSL—Cisco 857 Integrated Service Router with an ADSL interface.
- 876 ADSL—Cisco 876 Integrated Services Router with an ADSL interface.
- 877 ADSL—Cisco 877 Integrated Services Router with an ADSL interface.
- 878 888 G.SHDSL—Cisco 878 Integrated Services Router with a G.SHDSL interface.
- 1801 ADSLoPOTS—Cisco 1801 Integrated Services Router that provides ADSL over POTS.
- 1802 ADSLoISDN—Cisco 1802 Integrated Services Router that provides ADSL over ISDN.
- 1803 G.SHDSL—Cisco 1803 Integrated Services Router that provides 4-wire G.SHDSL.

Navigation Path

Go to the PVC Policy Page, page 62-55, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining ATM PVCs, page 62-51

Field Reference

Table 62-22 PVC Dialog Box

Element	Description
ATM Interface	The ATM interface on which the PVC is defined. Enter the name of an interface, subinterface, or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. **Note** We strongly recommend not defining an interface role that includes ATM interfaces from different interface cards. The different settings supported by each card type may cause deployment to fail.
Interface Card	The type of WAN interface card installed on the router or the router type. Supported card types are listed above. **Note** To ensure proper policy validation, we highly recommend that you define a value in this field. When you discover a live device, the correct interface card type will already be displayed. If you did not perform discovery on a live device, or if Security Manager cannot detect the type of interface card installed on the device, this field displays “Unknown”.
Settings tab	Defines basic PVC settings, such as the VPI/VCI and encapsulation. See PVC Dialog Box—Settings Tab, page 62-58.
QoS tab	Defines ATM traffic shaping and other quality-of-service settings for the PVC. See PVC Dialog Box—QoS Tab, page 62-61.

Table 62-22 PVC Dialog Box (continued)

Element	Description
Protocol tab	Defines the IP protocol mappings configured for the PVC (static maps or Inverse ARP). See PVC Dialog Box—Protocol Tab, page 62-64.
Advanced button	Defines F5 Operation, Administration, and Maintenance (OAM) settings for the PVC. See PVC Advanced Settings Dialog Box—OAM Tab, page 62-67.

PVC Dialog Box—Settings Tab

Use the Settings tab of the PVC dialog box to configure the basic settings of the PVC, including:

- ID settings.
- Encapsulation settings.
- Whether ILMI and Inverse ARP are enabled.
- The maximum number of PPPoE sessions.
- The static domain (VPN service) name to use for PPPoA.

Navigation Path

Go to the PVC Dialog Box, page 62-56, then click the **Settings** tab.

Related Topics

- PVC Dialog Box—QoS Tab, page 62-61
- PVC Dialog Box—Protocol Tab, page 62-64
- PVC Advanced Settings Dialog Box, page 62-66
- Defining ATM PVCs, page 62-51

Field Reference

Table 62-23 PVC Dialog Box—Settings Tab

Element	Description
PVC ID settings	
VPI	The virtual path identifier of the PVC. In conjunction with the VCI, identifies the next destination of a cell as it passes through a series of ATM switches on the way to its destination. Valid values for most platforms range from 0 to 255. For Cisco 2600 and 3600 Series routers using Inverse Multiplexing for ATM (IMA), valid values range from 0 to 15, 64 to 79, 128 to 143, and 192 to 207. **Note** VPI/VCI values must be unique for all the PVCs configured on a selected interface. VPI/VCI values are unique to a single link only and might change as cells traverse the ATM network.

Table 62-23 PVC Dialog Box—Settings Tab (continued)

Element	Description
VCI	The 16-bit virtual channel identifier of the PVC. In conjunction with the VPI, identifies the next destination of a cell as it passes through a series of ATM switches on the way to its destination. Valid values vary by platform. Typically, values up to 31 are reserved for special traffic (such as ILMI) and should not be used. 3 and 4 are invalid. **Note** VPI/VCI values must be unique for all the PVCs configured on a selected interface. VPI/VCI values are unique to a single link only and might change as cells traverse the ATM network.
Handle	An optional name to identify the PVC. The maximum length is 15 characters.
Management PVC (ILMI)	Does not apply when configuring the PVC on a subinterface. When selected, designates this PVC as the management PVC for this ATM interface by enabling communication with the Interim Local Management Interface (ILMI). ILMI is a protocol defined by the ATM Forum for setting and capturing physical layer, ATM layer, virtual path, and virtual circuit parameters on ATM interfaces. See Understanding ILMI, page 62-50. When deselected, this PVC does not act as the management PVC. This is the default. **Note** The VPI/VCI for the management PVC is typically set to 0/16.
Encapsulation settings	

Table 62-23 PVC Dialog Box—Settings Tab (continued)

Element	Description
Type	Does not apply when the Management PVC (ILMI) check box is enabled. The ATM adaptation layer (AAL) and encapsulation type to use on the PVC: • [blank]—The encapsulation type is not defined. (When deployed, aal5snap is applied.) • aal2—For PVCs dedicated to AAL2 Voice over ATM. AAL2 is used for variable bit rate (VBR) traffic, which can be either realtime (VBR-RT) or non-realtime (VBR-NRT). • aal5autoppp—Enables the router to distinguish between incoming PPP over ATM (PPPoA) and PPP over Ethernet (PPPoE) sessions and create virtual access for both PPP types based on demand. • aal5ciscoppp—For the proprietary Cisco version of PPP over ATM. • aal5mux—Enables you to dedicate the PVC to a single protocol, as defined in the Protocol field. • aal5nlpid—Enables ATM interfaces to work with High-Speed Serial Interfaces (HSSI) that are using an ATM data service unit (ADSU) and running ATM-Data Exchange Interface (DXI). • aal5snap—Supports Inverse ARP and incorporates the Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) that precedes the protocol datagram. This allows multiple protocols to traverse the same PVC.
Virtual Template	The virtual template used for PPP over ATM on this PVC. Enter the name of a virtual template interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. When a user dials in, the virtual template is used to configure a virtual access interface. When the user is done, the virtual access interface goes down and the resources are freed for other dial-in users. **Note** If you modify the virtual template settings on an existing PVC, you must enter the shutdown command followed by the no shutdown command on the ATM subinterface to restart the interface. This causes the newly configured parameters to take effect.

Table 62-23 PVC Dialog Box—Settings Tab (continued)

Element	Description
Protocol	Applies only when aal5mux is the defined encapsulation type. The protocol carried by the MUX-encapsulated PVC: • frame-relay—Frame-Relay-ATM Network Interworking (FRF.5) on the Cisco MC3810. • fr-atm-srv—Frame-Relay-ATM Service Interworking (FRF.8) on the Cisco MC3810. • ip—IP protocol. • ppp—IETF-compliant PPP over ATM. You must specify a virtual template when using this protocol type. • voice—Voice over ATM.
Additional settings	
Enable ILMI	When selected, enables ILMI management on this PVC. When deselected, ILMI management on this PVC is disabled.
Inverse ARP	When selected, the Inverse Address Resolution Protocol (Inverse ARP) is enabled on the PVC. When deselected, Inverse ARP is disabled. This is the default. Inverse ARP is used to learn the Layer 3 addresses at the remote ends of established connections. These addresses must be learned before the virtual circuit can be used. **Note** Use the Protocol tab to define static mappings of IP addresses instead of dynamically learning the addresses using Inverse ARP. See PVC Dialog Box—Protocol Tab, page 62-64.
PPPoE Max Sessions	The maximum number of PPP over Ethernet sessions that are permitted on the PVC.
VPN Service Name	The static domain name to use on this PVC. The maximum length is 128 characters. Use this option when you want PPP over ATM (PPPoA) sessions in the PVC to be forwarded according to the domain name supplied, without starting PPP.

PVC Dialog Box—QoS Tab

Use the QoS tab of the PVC dialog box to configure the ATM traffic shaping and other quality-of-service settings of the PVC, including:

- The limit on packets placed on transmission rings.
- The QoS service.
- Whether random detection is enabled.

These settings regulate the flow of traffic over the PVC by queuing traffic that exceeds the defined allowable bit rates.

Note QoS values are highly hardware dependent. Please refer to your router documentation for additional details about the settings that can be configured on your device.

Navigation Path

Go to the PVC Dialog Box, page 62-56, then click the **QoS** tab.

Related Topics

- PVC Dialog Box—Settings Tab, page 62-58
- PVC Dialog Box—Protocol Tab, page 62-64
- PVC Advanced Settings Dialog Box, page 62-66
- Defining ATM PVCs, page 62-51
- Quality of Service Policy Page, page 66-19
- Understanding Policing and Shaping Parameters, page 66-6

Field Reference

Table 62-24 PVC Dialog Box—QoS Tab

Element	Description
Tx Ring Limit	The maximum number of transmission packets that can be placed on a transmission ring on the WAN interface card (WIC) or interface. The range of valid values depends on the type of interface card selected in the Settings tab. See PVC Dialog Box—Settings Tab, page 62-58.
Traffic Shaping settings	
Traffic Shaping	The type of service to define on the PVC: • [null]—The bit rate is not defined. • ABR—Available Bit Rate. A best-effort service suitable for applications that do not require guarantees against cell loss or delays. • CBR—Constant Bit Rate service. Delay-sensitive data, such as voice or video, is sent at a fixed rate, providing a service similar to a leased line. • UBR—Unspecified Bit Rate service. A best-effort service suitable for applications that are tolerant to delay and do not require realtime responses. • UBR+—Unspecified Bit Rate service. Unlike UBR, UBR+ attempts to maintain a guaranteed minimum rate. • VBR-NRT—Variable Bit Rate - Non-Real Time service. A service suitable for non-realtime applications that are bursty in nature. VBR is more efficient than CBR and more reliable than UBR. • VBR-RT—Variable Bit Rate - Real Time service. A service suitable for realtime applications that are bursty in nature. For more information about each service class, see Understanding ATM Service Classes, page 62-48.

Table 62-24 PVC Dialog Box—QoS Tab (continued)

Element	Description
ABR	The following fields are displayed when ABR is selected as the Bit Rate: • PCR—The peak cell rate in kilobits per second (kbps). It specifies the maximum value of the ABR. • MCR—The minimum cell rate in kilobits per second (kbps). It specifies the minimum value of the ABR. The ABR varies between the MCR and the PCR. It is dynamically controlled using congestion control mechanisms.
CBR	The following field is displayed when CBR is selected as the Bit Rate: • Rate—The constant bit rate (also known as the average cell rate) for the PVC in kilobits per second (kbps). An ATM VC configured for CBR can send cells at this rate for as long as required.
UBR	The following field is displayed when UBR is selected as the Bit Rate: • PCR—The peak cell rate for output in kilobits per second (kbps). Cells in excess of the PCR may be discarded.
UBR+	The following fields are displayed when UBR+ is selected as the Bit Rate: • PCR—The peak cell rate for output in kilobits per second (kbps). Cells in excess of the PCR may be discarded. • MCR—The minimum guaranteed cell rate for output in kilobits per second (kbps). Traffic is always allowed to be sent at this rate. **Note** UBR+ requires Cisco IOS Software Release 12.4(2)XA or later, or version 12.4(6)T or later.
VBR-NRT	The following fields are displayed when VBR-NRT is selected as the Bit Rate: • PCR—The peak cell rate for output in kilobits per second (kbps). Cells in excess of the PCR may be discarded. • SCR—The sustained cell rate for output in kilobits per second (kbps). This value, which must be lower than or equal to the PCR, represents the maximum rate at which cells can be transmitted without incurring data loss. • MBS—The maximum burst cell size for output. This value represents the number of cells that can be transmitted above the SCR but below the PCR without penalty.

Table 62-24 PVC Dialog Box—QoS Tab (continued)

Element	Description
VBR-RT	The following fields are displayed when VBR-RT is selected as the Bit Rate: • Peak Rate—The peak information rate for realtime traffic in kilobits per second (kbps). • Average Rate—The average information rate for realtime traffic in kilobits per second (kbps). This value must be lower than or equal to the peak rate. • Burst—The burst size for realtime traffic, in number of cells. Configure this value if the PVC carries bursty traffic. These values configure traffic shaping between realtime traffic (such as voice and video) and data traffic to ensure that the carrier does not discard realtime traffic, for example, voice calls.
IP QoS settings	
Random Detect	When selected, enables Weighted Random Early Detection (WRED) or VIP-distributed WRED (DWRED) on the PVC. When deselected, WRED and DWRED are disabled. This is the default. WRED is a queue management method that selectively drops packets as the interface becomes congested. See Tail Drop vs. WRED, page 66-4.

PVC Dialog Box—Protocol Tab

Use the Protocol tab of the PVC dialog box to add, edit, or delete the protocol mappings configured for the PVC. You may configured static mappings or Inverse ARP (broadcast or nonbroadcast) for each PVC, but not both.

Note IP is the only protocol supported by Security Manager for protocol mapping on ATM networks. You cannot define protocol mappings on the Management PVC (ILMI).

Navigation Path

Go to the PVC Dialog Box, page 62-56, then click the **Protocol** tab.

Related Topics

- PVC Dialog Box—Settings Tab, page 62-58
- PVC Dialog Box—QoS Tab, page 62-61
- PVC Advanced Settings Dialog Box, page 62-66
- Defining ATM PVCs, page 62-51

Field Reference

Table 62-25 PVC Dialog Box—Protocol Tab

Element	Description
IP Protocol Mapping	Displays the IP protocol mappings configured for the PVC.
Add button	Opens the Define Mapping Dialog Box, page 62-65. From here you can define an IP protocol mapping.
Edit button	Opens the Define Mapping Dialog Box, page 62-65. From here you can edit the selected mapping.
Delete button	Deletes the selected mapping from the table.

Define Mapping Dialog Box

Use the Define Mapping dialog box to configure the IP protocol mappings to use on the ATM PVC. Mappings are required by the PVC to discover which IP address is reachable at the other end of a connection. Mappings can either be learned dynamically using Inverse ARP (InARP) or defined statically. Static mappings are best suited for simple networks that contain only a few nodes.

Note Inverse ARP is only supported for the aal5snap encapsulation type. See PVC Dialog Box—Settings Tab, page 62-58.

Tip Use the CLI or FlexConfigs to configure mappings for protocols other than IP.

Navigation Path

Go to the PVC Dialog Box—Protocol Tab, page 62-64, then click **Add** or **Edit**.

Related Topics

- PVC Dialog Box, page 62-56
- Defining ATM PVCs, page 62-51

Field Reference

Table 62-26 Define Mapping Dialog Box

Element	Description
IP Options	The type of IP protocol mapping to use: • IP Address—Select this option when using static mapping. Enter the address or the name of a network/host object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. • InARP—Inverse ARP. Select this option when using dynamic mapping. This allows the PVC to resolve its own network addresses without configuring a static map. Dynamic mappings age out and are refreshed periodically every 15 minutes by default. **Note** InARP can be used only when aal5snap is the defined encapsulation type for the PVC. See PVC Dialog Box—Settings Tab, page 62-58.
Broadcast Options	Indicates whether to use this map entry when sending IP broadcast packets (such as EIGRP updates): • Broadcast—The map entry is used for broadcast packets. • No Broadcast—The map entry is used only for unicast packets. • None—Broadcast options are disabled.

PVC Advanced Settings Dialog Box

Use the PVC Advanced Settings dialog box to configure F5 Operation, Administration, and Maintenance (OAM) functionality on an ATM PVC. OAM is used to detect connectivity failures at the ATM layer.

For more information, see Defining OAM Management on ATM PVCs, page 62-54.

Navigation Path

Go to the PVC Dialog Box, page 62-56, then click **Advanced**.

Related Topics

- PVC Policy Page, page 62-55

Field Reference

Table 62-27 PVC Advanced Settings Dialog Box

Element	Description
OAM tab	Defines loopback, connectivity check, and AIS/RDI settings. See PVC Advanced Settings Dialog Box—OAM Tab, page 62-67.
OAM-PVC tab	Enables OAM loopbacks and connectivity checks on the PVC. See PVC Advanced Settings Dialog Box—OAM-PVC Tab, page 62-69.

PVC Advanced Settings Dialog Box—OAM Tab

Use the OAM tab of the PVC Advanced Settings dialog box to define:

- The number of loopback cell responses that move the PVC to the down or up state.
- The number of alarm indication signal/remote defect indication (AIS/RDI) cells that move the PVC to the down or up state.
- The number and frequency of segment/end continuity check (CC) activation and deactivation requests that are sent on this PVC.

For more information, see Defining OAM Management on ATM PVCs, page 62-54.

Note The settings defined in this tab are dependent on the settings defined in the OAM-PVC tab. See PVC Advanced Settings Dialog Box—OAM-PVC Tab, page 62-69.

Navigation Path

Go to the PVC Advanced Settings Dialog Box, page 62-66, then click the **OAM** tab.

Related Topics

- PVC Dialog Box, page 62-56

Field Reference

Table 62-28 PVC Advanced Settings Dialog Box—OAM Tab

Element	Description
Retry settings	
Enable OAM Retry	When selected, OAM management settings can be defined. When deselected, OAM management settings cannot be defined. **Note** If Enable OAM Management is deselected in the OAM-PVC tab, these settings are saved in the device configuration but are not applied.
Down Count	The number of consecutive, unreceived, end-to-end loopback cell responses that cause the PVC to move to the down state. The default is 3.
Up Count	The number of consecutive end-to-end loopback cell responses that must be received in order to move the PVC to the up state. The default is 5.
Retry Frequency	The interval between loopback cell verification transmissions in seconds. The default is 1 second. If a PVC is up and a loopback cell response is not received within the specified interval (as defined in the Frequency field of the PVC-OAM tab), loopback cells are transmitted at the frequency defined here to verify whether the PVC is down. If the number of consecutive cells that do not receive a response matches the defined down count, the PVC is moved to the down state.
AIS-RDI settings	

Table 62-28 PVC Advanced Settings Dialog Box—OAM Tab (continued)

Element	Description
Enable AIS-RDI Detection	When selected, alarm indication signal (AIS) cells and remote defect indication (RDI) cells are used to report connectivity failures at the ATM layer of the PVC. When deselected, AIS/RDI cells are disabled. AIS cells notify downstream devices of the connectivity failure. The last ATM switch then generates RDI cells in the upstream direction towards the device that sent the original failure notification.
Down Count	The number of consecutive AIS/RDI cells that cause the PVC to go down. Valid values range from 1 to 60. The default is 1.
Up Count	The number of seconds after which a PVC is brought up if no AIS/RDI cells are received. Valid values range from 3 to 60 seconds. The default is 3.
Segment Continuity Check settings	
Enable Segment Continuity Check	When selected, OAM F5 continuity check (CC) activation and deactivation requests are sent to a device at the other end of a segment. When deselected, segment CC activation and deactivation requests are disabled. **Note** If Configure Continuity Check is deselected in the OAM-PVC tab, these settings are saved in the device configuration but are not applied.
Activation Count	The maximum number of times that the activation request is sent before the receipt of an acknowledgement. Valid values range from 3 to 600. The default is 3.
Deactivation Count	The maximum number of times that the deactivation request is sent before the receipt of an acknowledgement. Valid values range from 3 to 600. The default is 3.
Retry Frequency	The interval between activation/deactivation retries, in seconds. The default is 30 seconds.
End-to-End Continuity Check settings	
Enable End-to-End Continuity Check	When selected, OAM F5 continuity check (CC) activation and deactivation requests are sent to a device at the other end of the PVC. When deselected, segment CC activation and deactivation requests are disabled. **Note** If Configure Continuity Check is deselected in the OAM-PVC tab, these settings are saved in the device configuration but are not applied.
Activation Count	The maximum number of times that the activation request is sent before the receipt of an acknowledgement. Valid values range from 3 to 600. The default is 3.
Deactivation Count	The maximum number of times that the deactivation request is sent before the receipt of an acknowledgement. Valid values range from 3 to 600. The default is 3.

Table 62-28 PVC Advanced Settings Dialog Box—OAM Tab (continued)

Element	Description
Retry Frequency	The interval between activation/deactivation retries, in seconds. The default is 30 seconds.

PVC Advanced Settings Dialog Box—OAM-PVC Tab

Use the OAM-PVC tab of the PVC Advanced Settings dialog box to enable loopback cells and connectivity checks (CCs) on the PVC. These functions test the connectivity of the virtual connection.

For more information, see Defining OAM Management on ATM PVCs, page 62-54.

Note Use the OAM tab to define additional settings related to the settings on this tab. See PVC Advanced Settings Dialog Box—OAM Tab, page 62-67.

Navigation Path

Go to the PVC Advanced Settings Dialog Box, page 62-66, then click the **OAM-PVC** tab.

Related Topics

- PVC Dialog Box, page 62-56

Field Reference

Table 62-29 PVC Advanced Settings Dialog Box—OAM-PVC Tab

Element	Description
OAM settings	
Enable OAM Management	When selected, OAM loopback cell generation and OAM management are enabled on the PVC. When deselected, OAM loopback cells and OAM management are disabled. However, continuity checks can still be performed.
Frequency	The interval between loopback cell transmissions. Valid values range from 0 to 600 seconds.
Segment Continuity Check settings	
Segment Continuity Check	The current configuration of OAM F5 continuity checks performed on PVC segments: • None—Segment continuity checks (CC) are disabled. • Deny Activation Requests—The PVC rejects activation requests from peer devices, which prevents OAM F5 CC management from being activated on the PVC. • Configure Continuity Check—Segment CCs are enabled on the PVC. The router on which CC management is configured sends a CC activation request to the router at the other end of the segment, directing it to act as either a source or a sink. Segment CCs occur on a PVC segment between the router and a first-hop ATM switch.

Table 62-29 PVC Advanced Settings Dialog Box—OAM-PVC Tab (continued)

Element	Description
Direction	Applies only when CC management is enabled. The direction in which CC cells are transmitted: • both—CC cells are transmitted in both directions. • sink—CC cells are transmitted toward the router that initiated the CC activation request. • source—CC cells are transmitted away from the router that initiated the CC activation request.
Keep VC up after segment failure	When selected, the PVC is kept in the up state when CC cells detect connectivity failure. When deselected, the PVC is brought down when CC cells detect connectivity failure.
Keep VC up after end-to-end failure	When selected, specifies that if AIS/RDI cells are received, the PVC is not brought down because of end CC failure or loopback failure. When deselected, the PVC is brought down because of end CC failure or loopback failure.
End-to-End Continuity Check settings	
End-to-End Continuity Check	The current configuration of OAM F5 end-to-end continuity checks on the PVC: • None—End-to-end continuity checks (CC) are disabled. • Deny Activation Requests—The PVC rejects activation requests from peer devices, which prevents OAM F5 CC management from being activated on the PVC. • Configure Continuity Check—End-to-end CCs are enabled on the PVC. The router on which CC management is configured sends a CC activation request to the router at the other end of the connection, directing it to act as either a source or a sink. End-to-end CC monitoring is performed on the entire PVC between two ATM end stations.
Direction	Applies only when CC management is enabled. The direction in which CC cells are transmitted: • both—CC cells are transmitted in both directions. • sink—CC cells are transmitted toward the router that initiated the CC activation request. • source—CC cells are transmitted away from the router that initiated the CC activation request.
Keep VC up after end-to-end failure	When selected, the PVC is kept in the up state when CC cells detect connectivity failure. When deselected, the PVC is brought down when CC cells detect connectivity failure.

Table 62-29 PVC Advanced Settings Dialog Box—OAM-PVC Tab (continued)

Element	Description
Keep VC up after segment failure	When selected, specifies that if AIS/RDI cells are received, the PVC is not brought down because of a segment CC failure. When deselected, the PVC is brought down because of a segment CC failure.

PPP on Cisco IOS Routers

The Point-to-Point Protocol (PPP), as defined in RFC 1661, provides a method for transporting packets between two devices or hosts using physical or logical links. PPP is a Layer 2 data-link protocol that can work with multiple Layer 3 network-layer protocols, including IP, IPX, and AppleTalk.

PPP is used in many common scenarios, such as:

- Connecting remote users to a central network over dial-in connections.
- Connecting the gateway of an enterprise network to an ISP for internet access.
- Connecting two LANs (for example, a central office and a branch office) to exchange data between them.

PPP connectivity is established in stages:

1. First, a Link Control Protocol (LCP) establishes, configures, and tests the data-link connection.
2. (Optional) Authentication verifies the identity of the two parties.
3. A family of Network Control Protocols (NCPs) establishes and configures the necessary network-layer protocols.

The PPP policy in Security Manager provides a method for configuring selected parameters that are negotiated between the two nodes during the LCP stage, including authentication (typically CHAP or PAP) and Multilink PPP (MLP). For more information about MLP, see Defining Multilink PPP Bundles, page 62-75.

The following topics describe the tasks you perform to create PPP policies on Cisco IOS routers:

- Defining PPP Connections, page 62-72
- Defining Multilink PPP Bundles, page 62-75

Understanding Multilink PPP (MLP)

MLP, as defined in RFC 1990, is a method for splitting, recombining, and sequencing datagrams across multiple logical data links. MLP was originally designed to exploit multiple bearer channels in ISDN, but it can be used whenever multiple PPP links connect two systems, including asynchronous links.

MLP spreads inbound and outbound traffic across multiple physical WAN links (known collectively as a bundle) while providing the following benefits:

- Packet fragmentation and reassembly
- Proper sequencing
- Multivendor interoperability
- Load balancing

As shown in Figure 62-3, traffic routed across an MLP link is fragmented, with the fragments being sent across the different physical links. At the remote end of the link, the fragments are reassembled and forwarded to the next hop toward their ultimate destination. By using multiple physical links, MLP provides a way to temporarily use the additional bandwidth afforded by these links.

Figure 62-3 Multilink PPP

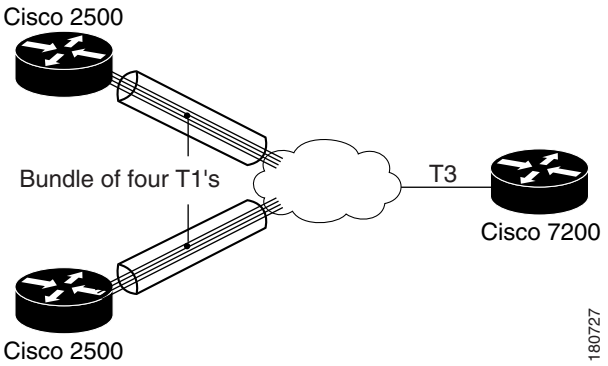

Every MLP bundle is controlled by a single interface, the bundle director, which is a virtual-access interface. This interface is created in the background when the bundle is first created. The physical interface becomes part of the bundle that is managed by the bundle director. Bundles are also used when you create a multilink group consisting of a multilink interface and its associated serial interfaces, which is a setup that is often found in static, leased-line environments.

MLP uses an endpoint discriminator to identify the system transmitting a packet. By default, this discriminator is based on the hostname of the router, but it can also be based on other criteria, such as the IP address or MAC address of the interface, a telephone number, or a user-defined string. If the endpoint discriminator matches the discriminator of an existing link, the new link is added to the matching bundle. If no match exists, a new bundle is created. When authentication is used, a new bundle is established whenever there is a mismatch in either the discriminator or the authentication information exchanged between the two nodes.

Related Topics

- Defining Multilink PPP Bundles, page 62-75
- PPP on Cisco IOS Routers, page 62-71

Defining PPP Connections

When you define a PPP connection, the first step is to select the interface on which PPP should be enabled. You must select one of the following interface types:

- Async
- Group-Async
- Serial
- High-Speed Serial Interface (HSSI)
- Dialer
- BRI, PRI (ISDN)
- Virtual template

- Multilink

You cannot define PPP connections on:

- Subinterfaces.
- Serial interfaces with Frame Relay encapsulation.
- Virtual template interfaces defined as Ethernet or tunnel types (serial is supported).

Note You cannot configure PPP on serial interfaces that are configured for Frame Relay encapsulation. See Defining Basic Router Interface Settings, page 62-4.

Note Deployment might fail if you define PPP on a virtual template that is also used in an 802.1x policy. See Defining 802.1x Policies, page 64-4.

You can select one or more authentication protocols and define when authentication should be performed.

In addition, you can configure the authentication and authorization methods to use when performing AAA on a remote security server. You can either define a default method list to use for all PPP connections on the device or define a customized method list that applies to a specific connection.

Before You Begin

- Make sure that the device contains an interface on which PPP can be configured. See Basic Interface Settings on Cisco IOS Routers, page 62-1.

Related Topics

- Defining Multilink PPP Bundles, page 62-75
- PPP on Cisco IOS Routers, page 62-71

Step 1 Do one of the following:

- (Device view) Select **Interfaces > Settings > PPP/MLP** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > PPP/MLP** from the Policy Type selector. Select an existing policy or create a new one.

The PPP/MLP page is displayed. See Table 62-30 on page 62-76 for a description of the fields on this page.

Step 2 Click the **Add** button beneath the table to display the PPP dialog box.

Step 3 In the Interface field, enter the name of the interface or interface role on which you want to define the PPP connection, or click **Select** to select an interface role from a list or to create a new one.

Step 4 (Optional) On the PPP tab, define authentication for the PPP connection:

a. Select one or more authentication protocols.

b. Select one or more authentication options. These options determine when to perform authentication (callin, callout, and callback), whether to use one-time passwords, and whether to allow a mobile station in a PDSN configuration to receive Simple IP and Mobile IP services without using CHAP or PAP.

Note The Call Back option only enables authentication during callback. Use the CLI or FlexConfigs to configure the callback feature on the device.

c. See PPP Dialog Box—PPP Tab, page 62-78 for a description of the fields on this tab.

Step 5 (Optional) When using a remote AAA server to perform authentication, select Default List or Custom Method List in the Authenticate Using field, then define the methods to use in the Prioritized Method List field.

Note If you modify the default list, your changes affect all PPP connections on the devices that use this list. If you leave this field blank, authentication is performed using the local database on the device.

Step 6 (Optional) When using a remote AAA server to perform authorization, select AAA Policy Default List or Custom Method List, then define the methods to use in the Prioritized Method List field.

Note If you choose AAA Policy Default List, the device uses the default authorization methods defined in the AAA policy. See Defining AAA Services, page 63-4.

Step 7 (Optional) Define the username and password to send in response to PAP authentication requests.

Note If you entered the encrypted version of the password, select the **Encrypted** check box.

Step 8 (Optional) Define a different hostname to send in all CHAP challenges and responses in place of the router's own hostname.

Note If you entered the encrypted version of the password, select the **Encrypted** check box.

Step 9 (Optional) To enable Multilink PPP on this connection, click the **MLP** tab. See Defining Multilink PPP Bundles, page 62-75.

Step 10 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the PPP table.

Note To edit a PPP connection, select it from the table, then click **Edit**. To remove a PPP connection, select it, then click **Delete**.

Step 11 Repeat Step 2 to Step 10 to define PPP connections on additional interfaces. Only one PPP connection may be defined on an interface.

Defining Multilink PPP Bundles

You enable Multilink PPP (MLP) on the selected interface by selecting the check box at the top of the Multilink tab in the PPP dialog box. You can optionally enable Multiclass Multilink PPP (MCMP), which prevents delay-sensitive traffic from fragmentation, and interleaving, which enables packets to be interspersed among the fragments of larger packets. If you want to restrict a serial interface to a specific bundle, you can select the multilink interface that represents that bundle.

In addition, you can optionally modify the following default settings:

- The maximum fragment delay.
- The endpoint discriminator that identifies the router when negotiating the use of MLP.
- The maximum receive reconstructed unit (MRRU) permitted by the router and its peers.
- The maximum queue depth for first-in, first-out (FIFO) and non-FIFO queues.

Before You Begin

- Select the interface on which the PPP connection should be enabled.

Related Topics

- Defining PPP Connections, page 62-72
- PPP on Cisco IOS Routers, page 62-71

Step 1 In the PPP dialog box, click the **MLP** tab. See PPP Dialog Box—MLP Tab, page 62-80 for a description of the fields on this tab.

Step 2 Select the **Enable Multilink Protocol (MLP)** check box.

Step 3 (Optional) Configure one or more of the following options:

a. Whether to enable the multiclass feature that prevents delay-sensitive traffic from being fragmented. This is achieved by placing delay-sensitive traffic in a separate class from regular traffic.

b. Whether to enable the interleaving of packets among the fragments of larger packets on the MLP bundle.

c. Whether to restrict the physical link to joining only a designated multilink-group (defined by selecting a multilink interface). If a peer at the other end of the link tries to join a different bundle, the connected is severed.

d. Whether to modify the default amount of time required to transmit a fragment on the MLP bundle. The default is 30 milliseconds.

Note If you enable interleaving without defining a fragment delay, the default delay of 30 seconds is configured. This value does not appear in Security Manager or in the device configuration.

Step 4 (Optional) Under Endpoint, modify the default endpoint discriminator used on the MLP bundle.

The endpoint discriminator is used to identify the router on the MLP bundle. The default endpoint discriminator is either the globally configured hostname, or the PAP username or CHAP hostname (depending on the authentication protocol being used), if you configured those values on the PPP tab. See Defining PPP Connections, page 62-72.

Step 5 (Optional) In the MRRU fields, modify the default maximum packet size that the router (local) or the peer (remote) is capable of receiving.

Step 6 (Optional) Modify the default maximum size of link transmit queues when using FIFO and non-FIFO (QoS) queuing.

Step 7 Click **OK** to close the dialog box. Your definitions are displayed on the PPP page.

PPP/MLP Policy Page

Use the PPP/MLP page to create, edit, and delete PPP connections on the router. For more information, see Defining PPP Connections, page 62-72.

Navigation Path

- (Device view) Select **Interfaces > Settings > PPP/MLP** from the Policy selector.
- (Policy view) Select **Router Interfaces > Settings > PPP/MLP** from the Policy Type selector. Right-click **PPP/MLP** to create a policy, or select an existing policy from the Shared Policies selector.

Related Topics

- PPP on Cisco IOS Routers, page 62-71
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 62-30 PPP/MLP Page

Element	Description
Interface	The interface that is configured for PPP/MLP.
Authentication	The types of authentication used on the PPP connection.
Authorization	The method list used for AAA authorization on the PPP connection.
Multilink	Indicates whether Multilink PPP (MLP) is enabled on this PPP connection.
Endpoint	The type of default endpoint discriminator to use when negotiating the use of MLP with the peer.
Multiclass	Indicates whether the Multiclass Multilink PPP (MCMP) feature is enabled on this PPP connection.
Group	The number of the multilink-group interface to which the physical link is restricted.
Interleave	Indicates whether the PPP multilink interleave feature is enabled on this PPP connection.
Add button	Opens the PPP Dialog Box, page 62-77. From here you can define the authentication and multilink settings for the PPP connection.
Edit button	Opens the PPP Dialog Box, page 62-77. From here you can edit the selected PPP connection.
Delete button	Deletes the selected PPP connection from the table.

PPP Dialog Box

Use the PPP dialog box to configure PPP connections on the router. When you configure a PPP connection, you can define the type of authentication and authorization to perform and define multilink parameters.

Navigation Path

Go to the PPP/MLP Policy Page, page 62-76, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining PPP Connections, page 62-72

Field Reference

Table 62-31 PPP Dialog Box

Element	Description
Interface	The interface on which PPP encapsulation is enabled. Enter the name of an interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. The following interface types support PPP: • Async • Group-Async • Serial • High-Speed Serial Interface (HSSI) • Dialer • BRI, PRI (ISDN) • Virtual template • Multilink You cannot define PPP on: • Subinterfaces. • Serial interfaces with Frame Relay encapsulation. • Virtual template interfaces defined as Ethernet or tunnel types (serial is supported). **Note** You can define only one PPP connection per interface. **Note** Deployment might fail if you define PPP on a virtual template that is also used in an 802.1x policy. See 802.1x Policy Page, page 64-5.
PPP tab	Defines the type of authentication and authorization to perform on the PPP connection. See PPP Dialog Box—PPP Tab, page 62-78.

Table 62-31 PPP Dialog Box (continued)

Element	Description
MLP tab	Defines how to split and recombine sequential datagrams across multiple logical data links using Multilink PPP (MLP). See PPP Dialog Box—MLP Tab, page 62-80. This tab is greyed out and cannot be opened for devices that do not support the configuration settings.

PPP Dialog Box—PPP Tab

Use the PPP tab of the PPP dialog box to define the types of authentication and authorization to perform on the PPP connection.

Navigation Path

Go to the PPP Dialog Box, page 62-77, then click the **PPP** tab.

Related Topics

- PPP Dialog Box—MLP Tab, page 62-80

Field Reference

Table 62-32 PPP Dialog Box—PPP Tab

Element	Description
Authentication settings	
PPP Encapsulation	When selected, indicates that PPP encapsulation is enabled for the selected interface. This field is read-only.
Protocol	The authentication protocols to use: • CHAP—Challenge-Handshake Authentication Protocol. • PAP—Password Authentication Protocol. • MS-CHAP—Version 1 of the Microsoft version of CHAP (RFC 2433). • MS-CHAP-2—Version 2 of the Microsoft version of CHAP (RFC 2759). • EAP—Extensible Authentication Protocol. You may select one or more authentication protocols, as required.

Table 62-32 PPP Dialog Box—PPP Tab (continued)

Element	Description
Options	The authentication options to use: • Call In—When selected, authentication is performed on incoming calls. • Call Out—When selected, authentication is performed on outgoing calls. • Call Back—When selected, authentication is performed on callback. • One Time—When selected, one-time passwords are used for authentication. One-time passwords are considered highly secure since each one is used only once. When deselected, one-time passwords are not used. **Note** AAA authentication must be enabled in order to use one-time passwords. See AAA Policy Page, page 63-6. One-time passwords cannot be used with CHAP. • Optional—When selected, allows a mobile station in a Packet Data Serving Node (PDSN) configuration to receive Simple IP and Mobile IP services without using CHAP or PAP. When deselected, mobile stations must use CHAP or PAP to receive Simple IP and Mobile IP services.
Authenticate Using	AAA authentication settings for the PPP connection: • PPP Default List—Defines a default list of methods to be queried when authenticating a user for PPP. Enter the names of one or more AAA server group objects (up to four) in the Prioritized Method List field, or click **Select** to select it. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authenticate users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Tip** After you create the default list for one PPP connection, you can use it for other PPP connections on this device. • Prioritized Method List—Defines a sequential list of methods to be queried when authenticating a user for this PPP connection only. **Note** Leave this field blank to perform authentication using the local database on the router.
PAP Authentication settings	
Username	The username to send in PAP authentication requests. The username is case sensitive.

Table 62-32 PPP Dialog Box—PPP Tab (continued)

Element	Description
Password	The password to send in PAP authentication requests. Enter the password again in the Confirm field. The password can contain 1 to 25 uppercase or lowercase alphanumeric characters. The password is case sensitive. The username and password are sent if the peer requests the router to authenticate itself using PAP.
Encrypted Password	When selected, this indicates that the password you entered is already encrypted. When deselected, this indicates that the password you entered is in clear text.
CHAP Authentication settings	
Hostname	By default, the router uses its hostname to identify itself to the peer. If required, you can enter a different hostname to use for all CHAP challenges and responses. For example, use this field to specify a common alias for all routers in a rotary group.
Secret	The secret used to compute the response value for any CHAP challenge from an unknown peer. Enter the secret again in the Confirm field.
Encrypted Secret	When selected, this indicates that the password you entered is already encrypted. When deselected, this indicates that the password you entered is in clear text.
Authorization settings	
Authorize Using	AAA authorization settings for the PPP connection: • AAA Policy Default List—Uses the default authorization method list that is defined in the device's AAA policy. See AAA Policy Page, page 63-6. • Prioritized Method List—Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select it. Use the up and down arrows to define the order in which selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** Leave this field blank to perform authorization using the local database on the router.

PPP Dialog Box—MLP Tab

Use the MLP tab of the PPP dialog box to define Multilink PPP (MLP) parameters for the selected PPP connection.

Navigation Path

Go to the PPP Dialog Box, page 62-77, then click the **MLP** tab.

Related Topics

- PPP Dialog Box—PPP Tab, page 62-78

Field Reference

Table 62-33 PPP Dialog Box—MLP Tab

Element	Description
Enable Multilink PPP (MLP)	When selected, MLP is enabled on this PPP connection. When deselected, MLP is disabled.
Allow Multiple Data Classes	When selected, enables multiple data classes on the MLP bundle. Delay-sensitive traffic is placed into Class 1, where it can be interleaved but never fragmented. Normal data traffic is placed into Class 0, which is subject to fragmentation just as regular multilink packets are. When deselected, all traffic is subject to fragmentation.
Enable Interleaving of Packets Among Fragments of Larger Packets	When selected, enables the interleaving of packets among the fragments of larger packets on the MLP bundle. **Note** If you enable interleaving without defining a fragment delay, the default delay of 30 seconds is configured. This value does not appear in Security Manager or in the device configuration. When deselected, interleaving is disabled. **Note** Serial interfaces do not support interleaving.
Multilink Group	Applies only to serial, Group-Async, and multilink interfaces. Restricts the physical link to the selected multilink-group interface. Enter the name of a multilink interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. This option is typically used in static leased-line environments, where the remote systems to which the device's serial lines are connected are known in advance. In effect, this option dedicates a specific interfaces to a particular user, even when that user is not connected. If a peer at the other end of the link tries to join a different bundle, the connected is severed.
Maximum Fragment Delay	The maximum amount of time that should be required to transmit a fragment on the MLP bundle. Valid values range from 1 to 1000 milliseconds. Fragment size is determined by the defined fragment delay and the bandwidth of the links. **Note** Serial interfaces do not support this feature.

Table 62-33 PPP Dialog Box—MLP Tab (continued)

Element	Description
Endpoint Type	The identifier used by the router when transmitting packets on the MLP bundle: • [null]—Negotiation is conducted without using an endpoint discriminator. (No CLI command is generated.) • Hostname—The hostname of the router. This option is useful when multiple routers are using the same username to authenticate but have different hostnames. • IP—A defined IP address. Enter the address or the name of a network/host object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. • MAC—The MAC address of a specific interface. Enter the name of the interface or interface role, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. • None—Negotiation is conducted without using an endpoint discriminator. (The relevant CLI command is generated, but no endpoint discriminator is provided.) This option is useful when the router is connected to a malfunctioning peer that does not handle the endpoint discriminator properly. • Phone—An E.164-compliant telephone number. Enter the number in the field displayed. • String—A character string. Enter the string in the field displayed. The default endpoint discriminator is either the globally configured hostname, or the PAP username or CHAP hostname (depending on the authentication protocol being used), if you have configured those values on the PPP tab.
MRRU Local Peer	The maximum receive reconstructed unit (MRRU) value of the local peer. This value represents the maximum size packet that the local router is capable of receiving. Valid values range from 128 to 16384 bytes. The default is the maximum transmission unit (MTU) of the multilink group interface and 1524 bytes for all other interfaces.
MRRU Remote Peer	The maximum receive reconstructed unit (MRRU) value of the remote peer. This value represents the maximum size packet that the remote peer is capable of receiving. Valid values range from 128 to 16384 bytes. The default is 1524 bytes.
Maximum FIFO Queue Size	The maximum queue depth when the bundle uses first-in, first-out (FIFO) queuing. Valid values range from 2 to 255 packets. The default is 8.
Maximum QoS Queue Size	The maximum queue depth when the bundle uses non-FIFO queuing. Valid values range from 2 to 255 packets. The default is 2.

Router Device Administration

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

This chapter contains the following topics:

- AAA on Cisco IOS Routers, page 63-2
- AAA Policy Page, page 63-6
- User Accounts and Device Credentials on Cisco IOS Routers, page 63-14
- Accounts and Credential s Policy Page, page 63-16
- Bridging on Cisco IOS Routers, page 63-18
- Bridging Policy Page, page 63-21
- Time Zone Settings on Cisco IOS Routers, page 63-22
- Clock Policy Page, page 63-23
- CPU Utilization Settings on Cisco IOS Routers, page 63-25
- CPU Policy Page, page 63-26
- HTTP and HTTPS on Cisco IOS Routers, page 63-28
- HTTP Policy Page, page 63-31
- Line Access on Cisco IOS Routers, page 63-35
- Console Policy Page, page 63-42
- VTY Policy Page, page 63-50
- Optional SSH Settings on Cisco IOS Routers, page 63-63
- Secure Shell Policy Page, page 63-64
- SNMP on Cisco IOS Routers, page 63-66
- SNMP Policy Page, page 63-69
- DNS on Cisco IOS Routers, page 63-74
- DNS Policy Page, page 63-76
- Hostnames and Domain Names on Cisco IOS Routers, page 63-77
- Hostname Policy Page, page 63-78
- Memory Settings on Cisco IOS Routers, page 63-78

- Memory Policy Page, page 63-79
- Secure Device Provisioning on Cisco IOS Routers, page 63-81
- Secure Device Provisioning Policy Page, page 63-85
- DHCP on Cisco IOS Routers, page 63-87
- DHCP Policy Page, page 63-92
- NTP on Cisco IOS Routers, page 63-96
- NTP Policy Page, page 63-98

AAA on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Authentication, authorization, and accounting (AAA) network security services provide the primary framework through which you set up access control on your Cisco IOS router. Use the AAA policy in Security Manager to enable AAA functionality on Cisco IOS routers and to configure default AAA settings. The default settings that you define in this policy can be used in other policies, such as HTTP and line access (console and VTY) policies. Enabling AAA functionality is a prerequisite for any device policy that makes use of AAA, including NAC, SDP, and 802.1x.

For more information about AAA, see:

- Supported Authorization Types, page 63-2
- Supported Accounting Types, page 63-3
- Understanding Method Lists, page 63-3

To configure a AAA policy, see:

- Defining AAA Services, page 63-4

Related Topics

- Understanding AAA Server and Server Group Objects, page 6-27
- Line Access on Cisco IOS Routers, page 63-35

Supported Authorization Types

AAA authorization enables you to limit the services available to an authenticated user. Security Manager supports the following types of authorization:

- Network—Authorizes various types of network connections, such as PPP, SLIP, and ARAP.
- EXEC—Authorizes the launching of EXEC (CLI) sessions.
- Command—Authorizes the use of all EXEC mode commands that are associated with specific privilege levels.

When authorization is enabled, the router uses information retrieved from the user's profile to configure the user session. The profiles are located either in the local user database or on a security server. Users are granted access to a requested service only if the profile allows it.

Related Topics

- Supported Accounting Types, page 63-3
- Understanding Method Lists, page 63-3
- Defining AAA Services, page 63-4
- AAA on Cisco IOS Routers, page 63-2

Supported Accounting Types

AAA accounting enables you to track the services the users are accessing and the amount of network resources that they are consuming. Security Manager supports the following accounting types:

- Connection—Records information about all outbound connections made from this device, such as Telnet, local-area transport (LAT), TN3270, packet assembler/disassembler (PAD), and rlogin connections.

  For example, a RADIUS connection accounting record for an outbound Telnet connection includes such information as the port and IP address of the network access server (NAS), the start and end times of the connection, the identity of the user, and the number of packets that were transmitted during the session.
- EXEC—Records information about user EXEC (CLI) sessions on the devices, including the username, date, start and stop times, and the IP address of the NAS. For dial-in users, the record includes the telephone number from which the call originated.
- Command—Records information about the EXEC commands executed on the device by users with specific privilege levels. Each command accounting record includes a list of the commands executed for that privilege level, the date and time each command was executed, and the name of the user who executed it.

For each accounting type, you can choose whether you want to generate an accounting record at the start and end of each user session or only at the end.

When AAA accounting is enabled, the router sends accounting records of user activity to the TACACS+ or RADIUS security server. Each accounting record contains accounting attribute-value (AV) pairs and is stored on the security server. This data can later be analyzed for network management, client billing, and auditing purposes.

Related Topics

- Supported Accounting Types, page 63-3
- Understanding Method Lists, page 63-3
- Defining AAA Services, page 63-4
- AAA on Cisco IOS Routers, page 63-2

Understanding Method Lists

A method list is a sequential list describing the methods to use to perform a particular AAA function. In Security Manager, you define method lists by selecting AAA server groups, which are reusable objects that typically contain one or more AAA servers running the same protocol, such as RADIUS or TACACS+. Method lists enable you to designate one or more security protocols to be used for each AAA function, thus ensuring a backup system if the initial method fails.

Note Security Manager also contains predefined AAA server group objects for using the enable password or a local database. See Predefined AAA Authentication Server Groups, page 6-30.

For each AAA function, the device initially uses the first method defined in the list. If that method fails to respond, the device selects the next method in the list. This process continues until there is successful communication with a listed method, or all methods defined in the method list are exhausted.

Note The device attempts to communicate with the next listed method only when there is no response from the previous method. If the AAA service fails at any point in this cycle—meaning that the security server or local username database responds by denying the user access or services—the process stops and no other methods are attempted.

Related Topics

- Supported Authorization Types, page 63-2
- Supported Accounting Types, page 63-3
- Defining AAA Services, page 63-4
- AAA on Cisco IOS Routers, page 63-2

Defining AAA Services

To define AAA services on a Cisco IOS router, you must first enable AAA functionality on the router. After you do this, you can define the kind of functionality (authentication, authorization, and accounting) that you want the device to implement. You must define a method list for each function, including lists for each type of authorization and accounting that you enable.

For example, if you want to configure EXEC authorization and command authorization, you must define one method list for EXEC authorization and other method lists for each privilege level on which command authorization is performed.

Note If you use RADIUS for authentication, you must use the same RADIUS server group for authorization as well.

Related Topics

- Understanding Method Lists, page 63-3
- AAA on Cisco IOS Routers, page 63-2
- Understanding AAA Server and Server Group Objects, page 6-27

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > AAA** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > AAA** from the Policy Type selector. Select an existing policy or create a new one.

The AAA page is displayed. See AAA Policy Page, page 63-6 for a description of the fields on this page.

Step 2 Define which login authentication methods to use on users who access the device:

a. On the Authentication tab (see AAA Page—Authentication Tab, page 63-6), select the **Enable Device Login Authentication** check box.

b. Enter the names of one or more AAA server group objects (up to four) in the Prioritized Method List field, or click **Select** select the object from a list or to create a new one. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used.

Note If you select None as a method, it must appear as the last method in the list.

Step 3 (Optional) In the Maximum Number of Attempts field, define the maximum number of unsuccessful authentication attempts to allow before a user is locked out.

Step 4 (Optional) Define which authorization methods to use on users who have been successfully authenticated:

a. Click the **Authorization** tab on the AAA page. See Table 63-3 on page 63-8 for a description of the fields on this tab.

b. Define method lists for one or more of the following types of authorization:

- Network
- EXEC
- Command—Click the **Add** button to display the Command Authorization dialog box (see Command Authorization Dialog Box, page 63-10). From here, you can select a privilege level and the method list to apply to it.

For more information about these authorization types, see Supported Authorization Types, page 63-2.

Note RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization.

Step 5 (Optional) Define which accounting methods to use on the activities performed by users:

a. Click the **Accounting** tab on the AAA page. See Table 63-5 on page 63-11 for a description of the fields on this tab.

b. Define method lists for one or more of the following types of accounting:

- Connection
- EXEC
- Command—Click the **Add** button to display the Command Accounting dialog box (see Command Accounting Dialog Box, page 63-13). From here, you can select a privilege level and the method list to apply to it.

For more information about these accounting types, see Supported Accounting Types, page 63-3.

c. For each accounting type defined above, select a value from the Accounting Process Notices list. This defines when to create an accounting record, at the beginning and end of the user process or only at the end.

d. For each accounting type defined above, select the **Enable broadcast to multiple servers** check box if you want accounting information sent simultaneously to the first server in each AAA server group defined in the method list.

AAA Policy Page

Use the AAA page to define the default authentication, authorization, and accounting methods to use on the router. You do this by configuring method lists, which define which methods to use and the sequence in which to use them.

Note You can use the method lists defined in this policy as default settings when you configure AAA on the router's console port and VTY lines. See Console Policy Page, page 63-42 and VTY Policy Page, page 63-50.

Navigation Path

- (Device view) Select **Platform > Device Admin > AAA** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > AAA** from the Policy Type selector. Right-click **AAA** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- AAA on Cisco IOS Routers, page 63-2
- Understanding AAA Server and Server Group Objects, page 6-27
- Console Policy Page, page 63-42
- VTY Policy Page, page 63-50

Field Reference

Table 63-1 AAA Page

Element	Description
Authentication tab	Defines the login authentication methods to use and the sequence in which to use them. See AAA Page—Authentication Tab, page 63-6.
Authorization tab	Defines the types of network, EXEC, and command authorization to perform and the methods to use for each type. See AAA Page—Authorization Tab, page 63-8.
Accounting tab	Defines types of connection, EXEC, and command accounting to perform and the methods to use for each type. See AAA Page—Accounting Tab, page 63-10.

AAA Page—Authentication Tab

Use the Authentication tab of the AAA page to define the methods used to authenticate users who access the device. Authentication methods are defined in a method list, which define the security protocols to use, such as LDAP, RADIUS, and TACACS+.

Note You can use the method list defined in this policy on the console and VTY lines that are used to communicate with the device. See Console Policy Page, page 63-42 and VTY Line Dialog Box—Authentication Tab, page 63-55.

Navigation Path

Go to the AAA Policy Page, page 63-6, then click the **Authentication** tab.

Related Topics

- Defining AAA Services, page 63-4
- Understanding Method Lists, page 63-3
- AAA Server Group Dialog Box, page 6-49
- Predefined AAA Authentication Server Groups, page 6-30

Field Reference

Table 63-2 AAA Page—Authentication Tab

Element	Description
Enable Device Login Authentication	When selected, enables the authentication of all users when they log in to the device, using the methods defined in the method list. When deselected, authentication is not performed.
Prioritized Method List	Defines a sequential list of methods to be queried when authenticating a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authenticate users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. Supported methods include Line, Local, Kerberos, LDAP, RADIUS, TACACS+, and None. **Note** If you select None as a method, it must appear as the last method in the list.
Maximum Number of Attempts	The maximum number of unsuccessful authentication attempts before a user is locked out. This feature is disabled by default. Valid values range from 1 to 65535. **Note** From the standpoint of the user, there is no distinction between a normal authentication failure and an authentication failure due to being locked out. The system administrator has to explicitly clear the status of a locked-out user using **clear** commands.

AAA Page—Authorization Tab

Use the Authorization tab of the AAA page to define the type of authorization services to enable on the device and the methods to use for each type. Security Manager supports the following types of authorization:

- Network—Authorizes various types of network connections, such as PPP.
- EXEC—Authorizes the launching of EXEC sessions.
- Command—Authorizes the use of all EXEC mode commands that are associated with specific privilege levels.

Note You can use the method lists defined in this policy on the console and VTY lines that are used to communicate with the device. See Console Policy Page, page 63-42 and VTY Line Dialog Box—Authentication Tab, page 63-55.

Navigation Path

Go to the AAA Policy Page, page 63-6, then click the **Authorization** tab.

Related Topics

- Defining AAA Services, page 63-4
- Supported Authorization Types, page 63-2
- Understanding Method Lists, page 63-3
- AAA Server Group Dialog Box, page 6-49
- Filtering Tables, page 1-48

Field Reference

Table 63-3 AAA Page—Authorization Tab

Element	Description
Network Authorization settings	
Enable Network Authorization	When selected, enables the authorization of network connections, such as PPP, SLIP, or ARAP connections, using the methods defined in the method list. When deselected, network authorization is not performed.

Table 63-3 AAA Page—Authorization Tab (continued)

Element	Description
Prioritized Method List	Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. Supported methods include LDAP, RADIUS, TACACS+, Local, and None. **Note** RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization. **Note** If you select None as a method, it must appear as the last method in the list.
EXEC Authorization settings	
Enable CLI/EXEC Operations Authorization	When selected, this type of authorization determines whether the user is permitted to open an EXEC (CLI) session, using the methods defined in the method list. When deselected, EXEC authorization is not performed.
Prioritized Method List	Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received.
Command Authorization settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Authorization Dialog Box, page 63-10. From here you can configure a command authorization definition.
Edit button	Opens the Command Authorization Dialog Box, page 63-10. From here you can edit the command authorization definition.
Delete button	Deletes the selected command authorization definitions from the table.

Command Authorization Dialog Box

Use the Command Authorization dialog box to define which methods to use when authorizing the EXEC commands that are associated with a given privilege level. This enables you to authorize all commands associated with a specific privilege level, from 0 to 15.

Navigation Path

From the AAA Page—Authorization Tab, page 63-8, click the **Add** button beneath the Command Authorization table.

Related Topics

- Defining AAA Services, page 63-4
- Supported Authorization Types, page 63-2
- Understanding Method Lists, page 63-3

Field Reference

Table 63-4 Command Authorization Dialog Box

Element	Description
Privilege Level	The privilege level for which you want to define a command accounting list. Valid values range from 0 to 15.
Prioritized Method List	Defines a sequential list of methods to be used when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. Supported methods include TACACS+, Local, and None. **Note** If you select None as a method, it must appear as the last method in the list.

AAA Page—Accounting Tab

Use the Accounting tab of the AAA page to define the type of accounting services to enable on the device and the methods to use for each type. Security Manager supports the following types of accounting:

- Connection—Records information about all outbound connections made from this device.
- EXEC—Records information about user EXEC sessions on the devices, including the username, date, start and stop times, and the IP address.
- Command—Records information about the EXEC commands executed on the device by users with specific privilege levels.

In addition, you use the Accounting page to determine when accounting records should be generated and whether they should be broadcast to more than one AAA server.

Note You can use the method lists defined in this policy on the console and VTY lines that are used to communicate with the device. See Console Policy Page, page 63-42 and VTY Line Dialog Box—Authentication Tab, page 63-55.

Navigation Path

Go to the AAA Policy Page, page 63-6, then click the **Accounting** tab.

Related Topics

- Defining AAA Services, page 63-4
- Supported Accounting Types, page 63-3
- Understanding Method Lists, page 63-3
- AAA Server Group Dialog Box, page 6-49
- Filtering Tables, page 1-48

Field Reference

Table 63-5 AAA Page—Accounting Tab

Element	Description
Connection Accounting settings	
Enable Connection Accounting	When selected, enables the recording of information about outbound connections (such as Telnet) made over this device, using the methods defined in the method list. When deselected, connection accounting is not performed.
Generate Accounting Records for	Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. • Stop Only—Generates an accounting record at the end of the user process only. • None—Disables this type of accounting.
Prioritized Method List	Defines a sequential list of methods to be queried when creating connection accounting records for a user. Enter the names of one or more AAA server group objects (up to 10 for IOS 12.4(22)T+, otherwise up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. Supported methods include LDAP, RADIUS, and TACACS+.

Table 63-5 AAA Page—Accounting Tab (continued)

Element	Description
Enable Broadcast to Multiple Servers	When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.
EXEC Accounting Settings	
Enable CLI/EXEC Operations Accounting	When selected, enables the recording of basic information about user EXEC sessions, using the methods defined in the method list. When deselected, EXEC accounting is not performed.
Generate Accounting Records for	See description Table 43-7 on page 43-17.
Prioritized Method List	Defines a sequential list of methods to be queried when creating connection accounting records for a user. Enter the names of one or more AAA server group objects (up to 10 for IOS 12.4(22)T+, otherwise up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it.
Enable Broadcast to Multiple Servers	When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group.
Command Accounting settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Generate Accounting Records for	The points in the process where the device sends an accounting notice to the accounting server.
Enable Broadcast	Whether accounting records are broadcast to multiple servers simultaneously.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Accounting Dialog Box, page 63-13. From here you can configure a command accounting definition.
Edit button	Opens the Command Accounting Dialog Box, page 63-13. From here you can edit the command accounting definition.
Delete button	Deletes the selected command accounting definitions from the table.

Command Accounting Dialog Box

Use the Command Accounting dialog box to define which methods to use when recording information about the EXEC commands that are executed for a given privilege level. Each accounting record includes a list of the commands executed for that privilege level, as well as the date and time each command was executed, and the name of the user who executed it.

Navigation Path

From the AAA Page—Accounting Tab, page 63-10, click the **Add** button beneath the Command Accounting table.

Related Topics

- Defining AAA Services, page 63-4
- Supported Accounting Types, page 63-3
- Understanding Method Lists, page 63-3

Field Reference

Table 63-6 Command Accounting Dialog Box

Element	Description
Privilege Level	The privilege level for which you want to define a command accounting list. Valid values range from 0 to 15.
Generate Accounting Records for	Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.
Prioritized Method List	Defines a sequential list of methods to be used when creating accounting records for a user. Enter the names of one or more AAA server group objects (up to 10 for IOS 12.4(22)T+, otherwise up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. TACACS+ is the only supported method, but you can select multiple AAA server groups configured with TACACS+. **Note** If you select None as a method, it must appear as the last method in the list.

Table 63-6 Command Accounting Dialog Box (continued)

Element	Description
Enable Broadcast to Multiple Servers	When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.

User Accounts and Device Credentials on Cisco IOS Routers

Accounts and credential policies define the contact information for accessing the router, including the privilege level provided to each user account. You can configure as many user accounts as required. However, the user account that Security Manager uses to connect to the router is always the one configured in the Device Properties page.

Additionally, you use device access policies to define the enable or enable secret password required to access privileged EXEC mode. This is the mode required to make any configuration changes on the router.

Note If you use this policy to define a password, be careful later not to unassign this policy without assigning a replacement policy before your next deployment. If you deploy a device access policy that removes this password and the device contains a different type of password not known to Security Manager, such as a line console password, you will not be able to configure this device in the future. This is because the device reverts to this unknown password if Security Manager removes the enable password that it had previously configured.

Related Topics

- Defining Accounts and Credential Policies, page 63-14

Defining Accounts and Credential Policies

This procedure describes how to define a device access policy on a Cisco IOS router. If the username that you configured on the Device Properties page to connect to the router (see Viewing or Changing Device Properties, page 3-41) matches one of the user accounts you defined in this policy, Security Manager updates the device credentials according to your policy definition.

If you change the password for the user defined in the device properties, which Security Manager uses to deploy configurations to the device, or change the enable password, Security Manager uses the existing credentials defined in the device properties to log into the device and deploy changes. After successful deployment, the device properties are then changed to use your new settings. For more information on credentials in device properties, see Device Credentials Page, page 3-45.

Note You can discover encrypted passwords, but any password you enter must be in clear text. If you discover an encrypted password and then modify it, the password is saved as clear text.

Related Topics

- User Accounts and Device Credentials on Cisco IOS Routers, page 63-14

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Accounts and Credentials** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Accounts and Credentials** from the Policy Type selector. Select an existing policy or create a new one.

The Accounts and Credentials page is displayed. See Table 63-7 on page 63-16 for a description of the fields on this page.

Step 2 Enter the password for switching to privileged EXEC mode on the router:

a. Select **Enable Password** or **Enable Secret Password**. The Enable Secret Password option offers better security than the Enable Password option by storing the password using MD5 encryption. This option is useful in environments in which the password crosses the network or is stored on a TFTP server.

Note After you set an enable secret password, you can switch to an enable password only if the enable secret is disabled or an older version of Cisco IOS software is being used, such as when running an older rxboot image.

b. Enter a password, then enter it again in the Confirm field. The password that you enter must be in clear text. If you are configuring the enable secret password, the password is encrypted on deployment.

Step 3 (Optional) Select the **Enable Password Encryption Service** check box to encrypt all passwords on the device. This includes, for example, the enable password, username passwords, authentication key passwords, console and VTY line access passwords, and BGP neighbor passwords.

We recommend using this feature to help prevent unauthorized individuals from viewing the passwords in your configuration file.

Note This option does not provide a high level of security and should not be used as a substitute for additional network security measures.

Step 4 To define new user accounts for the router:

a. Click the **Add** button under the table to display the User Accounts dialog box.

b. Enter the details for the new user. See Table 63-8 on page 63-18 for a description of the available fields.

c. Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the User Accounts table.

Note To edit a user account, select it from the User Accounts table, then click **Edit**. To remove a user account, select it, then click **Delete**.

Caution While deleting a user account, Cisco Security Manager times out and deployment fails. To avoid this, you can set up Security Manager to download despite an error. Enable **Allow Download on Error** in **Tools > Administrator > Deployments** in the Configuration Manager.

Accounts and Credential s Policy Page

Use the Accounts and Credentials page to define the enable password or enable secret password assigned to the router. In addition, you can define a list of usernames that can be used to access the router.

For more information, see Defining Accounts and Credential Policies, page 63-14.

Navigation Path

- (Device view) Select **Platform > Device Admin > Accounts and Credentials** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Accounts and Credentials** from the Policy Type selector. Right-click **Accounts and Credentials** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- User Accounts and Device Credentials on Cisco IOS Routers, page 63-14
- User Account Dialog Box, page 63-17
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-7 Accounts and Credentials Page

Element	Description
Enable Secret Password	The enable secret password for entering privileged EXEC mode on the router. This option offers better security than the Enable Password option. The enable secret password can contain between 1-25 alphanumeric characters. The first character must be a letter. Spaces are allowed, but leading spaces are ignored. Question marks are also allowed. **Note** You can discover an encrypted password, but any password you enter must be in clear text. If you modify an encrypted password, it is saved as clear text. **Note** After you set an enable secret password, you can switch to an enable password only if the enable secret is disabled or an older version of Cisco IOS software is being used, such as when running an older rxboot image.

Table 63-7 Accounts and Credentials Page (continued)

Element	Description
Enable Password	The enable password for entering privileged EXEC mode on the router. The enable password can contain between 1-25 alphanumeric characters. The first character must be a letter. Spaces are allowed, but leading spaces are ignored. Question marks are also allowed. **Note** You must enter the password in clear text.
Enable Password Encryption Service	When selected, encrypts all passwords on the device, including the enable password (which is otherwise saved in clear text). For example, use this option to encrypt username passwords, authentication key passwords, console and VTY line access passwords, and BGP neighbor passwords. This command is primarily used for keeping unauthorized individuals from viewing your passwords in your configuration file. When deselected, device passwords are stored unencrypted in the configuration file. **Note** This option does not provide a high level of network security. You should also take additional network security measures.
User Accounts Table	
Username	The username that can be used to access the router. The username must be a single word up to 64 characters in length. Spaces and quotation marks are not allowed.
Encryption	Indicates whether password information for the user is encrypted using MD5 encryption.
Privilege Level	The privilege level assigned to the user.
Add button	Opens the User Account Dialog Box, page 63-17. From here you can define a user account.
Edit button	Opens the User Account Dialog Box, page 63-17. From here you can edit the selected user.
Delete button	Deletes the selected user accounts from the table.

User Account Dialog Box

Employ the User Account dialog box to define a username and password combination that can be used by Security Manager to access the router. You can also define the privilege level of the user account, which determines whether you can configure all commands on this router or only a subset of them.

Note Remember—there may be additional user accounts defined on the router using other methods, such as the CLI.

Navigation Path

Go to the Accounts and Credential s Policy Page, page 63-16, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining Accounts and Credential Policies, page 63-14
- User Accounts and Device Credentials on Cisco IOS Routers, page 63-14
- Understanding FlexConfig Policies and Policy Objects, page 7-2

Field Reference

Table 63-8 User Account Dialog Box

Element	Description
Username	The username for accessing the router.
Password	The password for accessing the router with this user account. **Note** You can discover an encrypted password, but any password you enter must be in clear text.
Confirm	Confirms the password for this user account.
Encrypt password using MD5	When selected, uses MD5 encryption to encrypt the password for this user account. This is the default. When deselected, the password is sent to the router unencrypted.
Privilege Level	The privilege level assigned to the user account. Valid values range from 0 to 15: • 0—Grants access to these commands only: **disable**, **enable**, **exit**, **help**, and **logout**. • 1—Enables nonprivileged access to the router (normal EXEC-mode use privileges). • 15—Enables privileged access to the router (traditional enable privileges). **Note** Levels 2-14 are not normally used in a default configuration, but custom configurations can be created by moving commands that are normally at level 15 to a lower level and commands that are normally at level 1 to a higher level. You can configure the privilege levels of commands using the CLI or by defining a FlexConfig.

Bridging on Cisco IOS Routers

Bridging policies enable you to perform transparent bridging (as specified in RFC 1286) on selected interfaces that you have configured to function as a bridge group. Security Manager supports integrated routing and bridging, which makes it possible to route a specific protocol between routed interfaces and bridge groups, or route a specific protocol between bridge groups. Local or unroutable traffic can be bridged among the bridged interfaces in the same bridge group, while routable traffic can be routed to other routed interfaces or bridge groups, as shown in Figure 63-1.

Using integrated routing and bridging, you can:

- Switch packets from a bridged interface to a routed interface.
- Switch packets from a routed interface to a bridged interface.

- Switch packets within the same bridge group.

Figure 63-1 Transparent Bridging

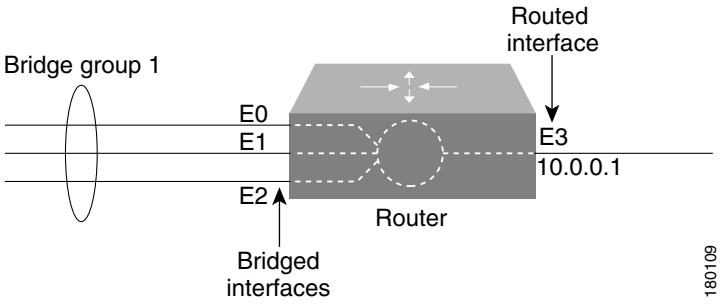

Related Topics

- Defining Bridge Groups, page 63-20
- Bridge-Group Virtual Interfaces, page 63-19

Bridge-Group Virtual Interfaces

Because bridging takes places at the data link layer and routing takes place at the network layer, they have different protocol configuration models. With IP, for example, bridge group interfaces belong to the same network and have a collective IP network address. In contrast, each routed interface represents a distinct network and has its own IP network address. Integrated routing and bridging uses the concept of a bridge-group virtual interface (BVI) to enable these interfaces to exchange packets for a given protocol. As shown in Figure 63-2, the interface number assigned to the BVI corresponds to the bridge group that the BVI represents. This number serves as the link between the virtual interface and the bridge group.

Figure 63-2 Bridge-Group Virtual Interface

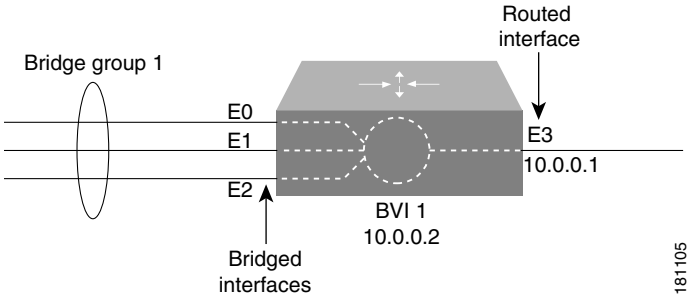

When you enable routing for a given protocol on the BVI, packets coming from a routed interface that are destined for a host in a bridged domain are routed to the BVI and then forwarded to the corresponding bridged interface. All traffic routed to the BVI is forwarded to the corresponding bridge group as bridged traffic. All routable traffic received on a bridged interface is routed to other routed interfaces as if it is coming directly from the BVI.

Note BVI interfaces are configured using the Interfaces policy. See Defining Basic Router Interface Settings, page 62-4. The BVI interface must have a corresponding bridge group with the same number; otherwise, deployment will fail.

Note When the bridge group contains more than two interfaces, add a BVI interface to the group to help prevent unicast flooding, which is a potential security issue.

Related Topics

- Defining Bridge Groups, page 63-20
- Bridging on Cisco IOS Routers, page 63-18

Defining Bridge Groups

You define a bridge group by selecting the L3 interfaces that are part of the bridge group and assigning the group a number. All bridge groups in Security Manager perform integrated routing and bridging on IP traffic only and use the standard Spanning Tree Protocol (IEEE 802.1D).

Note Use CLI commands or FlexConfigs to bridge other protocols, such as AppleTalk or IPX, and to use other spanning tree protocols, such as VLAN-Bridge. Concurrent routing and bridging is not supported.

Related Topics

- Bridging on Cisco IOS Routers, page 63-18
- Bridge-Group Virtual Interfaces, page 63-19

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Bridging** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Bridging** from the Policy Type selector. Select an existing policy or create a new one.

The Bridging page is displayed. See Table 63-9 on page 63-21 for a description of the fields on this page.

Step 2 Click the **Add** button under the table to display the Bridge Group dialog box. See Table 63-10 on page 63-22 for a description of the fields in this dialog box. From here you can define a bridge group.

Step 3 Enter a number to identify the bridge group.

Step 4 Enter the names of the interfaces and interface roles that are part of the bridge group, or click **Select** to select an interface role or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

You can select most Layer 3 interfaces, except X.25 and Integrated Services Digital Network (ISDN) bridged interfaces and certain types of logical interfaces (such as loopback, tunnel, null, and BVI). Each interface can be included in only one bridge group.

You can select a LAN subinterface only if the parent interface is configured with Inter-Switch Link (ISL) or 802.1Q encapsulation.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box. The bridge group is displayed in the table on the Bridging page.

Note To edit a bridge group, select it from the Groups table, then click **Edit**. To remove a bridge group, select it, then click **Delete**.

Bridging Policy Page

Use the Bridging page to define bridge groups that can perform integrated routing and bridging on the router. For more information, see Defining Bridge Groups, page 63-20.

Navigation Path

- (Device view) Select **Platform > Device Admin > Bridging** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Bridging** from the Policy Type selector. Right-click **Bridging** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Bridging on Cisco IOS Routers, page 63-18
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-9 Bridging Page

Element	Description
Group Number	The number that identifies the bridge group.
Group Interfaces	The interfaces and interface roles that are included in the bridge group.
Add button	Opens the Bridge Group Dialog Box, page 63-21. From here you can define a bridge group.
Edit button	Opens the Bridge Group Dialog Box, page 63-21. From here you can edit the bridge group.
Delete button	Deletes the selected bridge groups from the table.

Bridge Group Dialog Box

Use the Bridge Group dialog box to define bridge groups on the router. Each bridge group can contain multiple Layer 3 interfaces of various types, including serial interfaces.

Note All bridge groups use the standard Spanning Tree Protocol (IEEE 802.1D). Use CLI commands or FlexConfigs to bridge other protocols, such as AppleTalk or IPX, and to use other spanning tree protocols, such as VLAN-Bridge.

Navigation Path

Go to the Bridging Policy Page, page 63-21, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining Bridge Groups, page 63-20
- Bridging on Cisco IOS Routers, page 63-18
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 63-10 Bridge Group Dialog Box

Element	Description
Group Number	The number assigned to the bridge group. Valid values range from 1 to 255.
Group Interfaces	The interfaces that are included in the bridge group. Enter the name of one or more interfaces and interface roles, or click **Select** to select them. If the object that you want is not listed, click the **Create** button to create it. You can select most Layer 3 interfaces, including serial interfaces, provided the serial interface is configured with high-level data link control (HDLC) or Frame Relay encapsulation. Each interface can belong to only one bridge group. You can select a LAN subinterface only if the parent interface is configured with Inter-Switch Link (ISL) or 802.1Q encapsulation. **Note** Certain types of interfaces, such as loopback, tunnel, null, and BVI, cannot be bridged. **Note** Make sure that your bridge group does not prevent Security Manager from communicating with the device.

Time Zone Settings on Cisco IOS Routers

The local time on a Cisco IOS router is typically set using the clock set command in the CLI command or by dynamically deriving the time from an NTP server. You can adjust these time settings by defining the time zone in which the router resides and the start and end dates of Daylight Saving Time (DST) in that time zone.

Related Topics

- Defining Time Zone and DST Settings, page 63-22
- NTP on Cisco IOS Routers, page 63-96

Defining Time Zone and DST Settings

Security Manager enables you to define the time zone in which a Cisco IOS router is located. You can also define the start and end dates for Daylight Saving Time (DST).

Related Topics

- Defining NTP Servers, page 63-97
- Time Zone Settings on Cisco IOS Routers, page 63-22

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Clock** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Clock** from the Policy Type selector. Select an existing policy or create a new one.

The Clock page is displayed. See Table 63-11 on page 63-24 for a description of the fields on this page.

Step 2 Select the time zone in which the router is located. Time zones are listed according the number of hours behind or ahead of Greenwich Mean Time (GMT).

Step 3 (Optional) Select the method for determining the start and end dates for DST:

- Set by Date—Select this option when DST starts and ends on fixed dates. Continue with Step 4.
- Set by Day—Select this option when DST starts and ends on days whose specific dates vary from year to year. Continue with Step 5.
- None—Select this option when DST is not used.

Step 4 (When Set by Date is selected) Define the fixed dates when DST starts and ends:

a. Under Start, click the calendar icon, then click the appropriate date.

b. Select the hour and minute from the displayed lists.

c. Repeat steps a and b to configure the end date and time.

Step 5 (When Set by Day is selected) Select the **Specify Recurring Time** check box if you want to define a DST period *other* than the default, which is the period used throughout most of the United States.

Step 6 (When Specify Recurring Time is selected) Define the start and end of DST:

a. Under Start, select the month when DST begins.

b. Select the week of the month (1, 2, 3, 4, first, or last).

c. Select the day of the week.

d. Select the hour and minute from the displayed lists. For example, if DST begins at 1:00 a.m. on the last Sunday of each March, select March, last, Sunday, 1, and 00.

e. Repeat Steps a through d to configure the end date and time.

Clock Policy Page

Use the Clock page to configure the time zone in which the router is located and the settings for Daylight Saving Time (DST). For more information, see Time Zone Settings on Cisco IOS Routers, page 63-22.

Tip You can configure the local time on the router by defining an NTP policy or by configuring the **clock set** command using the CLI.

Navigation Path

- (Device view) Select **Platform > Device Admin > Clock** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Clock** from the Policy Type selector. Right-click **Clock** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- NTP Policy Page, page 63-98

Field Reference

Table 63-11 Clock Page

Element	Description
Device Time Zone	The time zone in which the router is located, expressed in relation to GMT (Greenwich Mean Time), also known as UTC (Coordinated Universal Time). **Caution** If you want to use the Command Line Interface (CLI) to configure the time zone on the router, you must use the required Time Zone acronym provided in the *Cisco IOS Configuration Fundamentals Command Reference* document. If you use any other format for the time zone and then use Security Manager to discover the router, Security Manager will not discover the time zone CLI.
Daylight Savings Time (Summer Time)	The type of DST to apply to the local time on the router: • Set by Date—Enables you to define the exact date and time when DST begins and ends. • Set by Day—Enables you to define the relative recurring date and time when DST begins and ends. For example, you can use this option when DST begins the last Sunday of March and ends the last Sunday of October. • None—Daylight savings time is not used.
Additional Set by Date fields	
Start	The date and time when DST begins: • Date—Click the calendar icon to select the start date. • Hour—Select the start hour. • Minute—Select the start minute.
End	The date and time when DST ends: • Date—Click the calendar icon to select the end date. • Hour—Select the end hour. • Minute—Select the end minute. **Note** Cisco IOS Software supports dates up to and including December 31st, 2035.
Additional Set by Day fields	

Table 63-11 Clock Page (continued)

Element	Description
Specify Recurring Time	When selected, the router implements DST according to the dates and times specified in this policy. When deselected, the router implements DST according to the schedule used throughout most of the United States.
Start	The relative date and time when daylight savings time begins: • Month—Select the month. • Week—Select the week of the month (1, 2, 3, 4, first, or last). • Weekday—Select the day of the week. • Hour—Select the hour. • Minute—Select the minute. For example, if DST begins at 1:00 a.m. on the last Sunday of each March, select March, last, Sunday, 1, and 00.
End	The relative date and time when daylight savings time ends: • Month—Select the month. • Week—Select the week of the month (1, 2, 3, 4, first, or last). • Weekday—Select the day of the week. • Hour—Select the hour. • Minute—Select the minute.

CPU Utilization Settings on Cisco IOS Routers

The CPU policy configures settings relating to CPU utilization. This policy provides you with methods for monitoring CPU resources and tracking processes that exceed a predetermined level of utilization.

Note The CPU policy is supported on routers running Cisco IOS Software Release 12.3(14)T or later.

Related Topics

- Defining CPU Utilization Settings, page 63-25

Defining CPU Utilization Settings

You can use Security Manager to modify the following default CPU utilization settings:

- The size of the CPU history table.
- The size of the extended CPU load history table.
- Whether to enable the automatic CPU Hog profiling.

In addition, you can optionally define:

- The CPU utilization level that causes a process to be included in the history table.

- The types of CPU utilization thresholds to enable. For each type of threshold, you can determine the threshold values that trigger notifications.

Related Topics

- CPU Utilization Settings on Cisco IOS Routers, page 63-25
- Logging on Cisco IOS Routers, page 65-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > CPU** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > CPU** from the Policy Type selector. Select an existing policy or create a new one.

The CPU page is displayed.

Step 2 (Optional) Define the CPU utilization settings of the router, as required. See Table 63-12 on page 63-26 for a description of the available fields.

CPU Policy Page

Use the CPU page to configure settings related to router CPU utilization, including the thresholds for sending log messages, the size of the CPU history table, and whether to enable automatic CPU Hog profiling.

For more information, see Defining CPU Utilization Settings, page 63-25.

Navigation Path

- (Device view) Select **Platform > Device Access > CPU** from the Policy selector.
- (Policy view) Select **Router Platform > Device Access > CPU** from the Policy Type selector. Right-click **CPU** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Memory Policy Page, page 63-79
- Syslog Logging Setup Policy Page, page 65-7
- Syslog Servers Policy Page, page 65-10

Field Reference

Table 63-12 CPU Page

Element	Description
CPU Utilization Statistics	Settings related to the history table for CPU utilization statistics: • History Table Entry Limit—The percentage of CPU utilization that a process must use to be included in the history table. • History Table Size—The length of time for which CPU statistics are stored in the history table. Valid values range from 5 to 86400 seconds (24 hours). The default is 600 seconds (10 minutes).

Table 63-12 CPU Page (continued)

Element	Description
CPU Total Utilization	The thresholds for total CPU utilization that trigger notifications: • Enable CPU Total Utilization—When selected, CPU total utilization thresholds are enabled. When deselected, these thresholds are disabled and do not trigger notifications. This is the default. • Maximum Total Utilization Resources—The percentage of CPU resources that, when usage *exceeds* this level for the defined interval, triggers a notification. • Maximum Total Utilization Violation Duration—The violation interval that triggers a maximum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours). • Minimum Total Utilization Resources—The percentage of CPU resources that, when usage *falls below* this level for the defined interval, triggers a notification. • Minimum Total Utilization Violation Duration—The violation interval that triggers a minimum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours).
CPU Interrupt Utilization	The thresholds for CPU interrupt utilization that trigger notifications: • Enable CPU Interrupt Utilization—When selected, CPU interrupt utilization thresholds are enabled. When deselected, these thresholds are disabled and do not trigger notifications. This is the default. • Maximum Interrupt Utilization Resources—The percentage of CPU resources that, when usage *exceeds* this level for the defined interval, triggers a notification. • Maximum Interrupt Utilization Violation Duration—The violation interval that triggers a maximum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours). • Minimum Interrupt Utilization Resources—The percentage of CPU resources that, when usage *falls below* this level for the defined interval, triggers a notification. • Minimum Interrupt Utilization Violation Duration—The violation interval that triggers a minimum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours).

Table 63-12 CPU Page (continued)

Element	Description
CPU Process Utilization	The thresholds for CPU process utilization that trigger notifications: • Enable CPU Process Utilization—When selected, CPU process utilization thresholds are enabled. When deselected, these thresholds are disabled and do not trigger notifications. This is the default. • Maximum Process Utilization Resources—The percentage of CPU resources that, when usage *exceeds* this level for the defined interval, triggers a notification. • Maximum Process Utilization Violation Duration—The violation interval that triggers a maximum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours). • Minimum Process Utilization Resources—The percentage of CPU resources that, when usage *falls below* this level for the defined interval, triggers a notification. • Minimum Process Utilization Violation Duration—The violation interval that triggers a minimum CPU threshold notification. Valid values range from 5 to 86400 seconds (24 hours).
Extended CPU History Size	The size of the history to collect for the extended CPU load, in increments of 5 seconds. Valid values range from 2 to 720. The default is 12, which is equivalent to a 1-minute history.
Enable Automatic CPU Hog Profiling	When selected, automatic CPU Hog profiling is enabled. This is the default. When deselected, automatic CPU Hog profiling is disabled. This feature predicts when a process could hog the CPU and begins profiling that process. **Note** To view the CPU Hog profile data, use the **show processes cpu autoprofile hog** command in the CLI.

HTTP and HTTPS on Cisco IOS Routers

Security Manager enables you to configure HTTP and HTTP over Secure Socket Layer (known as HTTP over SSL or HTTPS) server functionality on Cisco IOS routers. This feature provides SSL version 3.0 support for the HTTP 1.1 server.

A secure HTTP connection means that data sent to and received from an HTTP server are encrypted before being sent out over the internet. HTTP with SSL encryption provides a secure connection to allow such functions as configuring a router from a web browser.

In addition to providing access to the device via the Cisco web browser user interface, HTTP and HTTPS are used by device management applications, such as the Cisco Router and Security Device Manager (SDM), to communicate with the device.

Related Topics

- Defining HTTP Policies, page 63-29

Defining HTTP Policies

When you define an HTTP policy, you can:

- Enable and disable HTTP and SSL functionality on the router.
- Specify the ports used by each protocol.
- Optionally define a standard, numbered ACL that restricts access to the device using these protocols.

In addition, you can define the methods of AAA authentication and authorization methods to perform on users.

You must use caution when defining an HTTP policy, as your settings may affect communication between Security Manager (as well as other management applications that use these protocols) and the device.

Note As a general rule, Cisco IOS routers that have been discovered by Security Manager already have HTTPS enabled because Security Manager uses SSL as the default protocol for communicating with them. See Setting Up SSL on Cisco IOS Routers, page 2-4.

Before You Begin

- Enable AAA services on the router. See Defining AAA Services, page 63-4.

Related Topics

- HTTP and HTTPS on Cisco IOS Routers, page 63-28

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > HTTP** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Device Admin > Device Access > HTTP** from the Policy Type selector. Select an existing policy or create a new one.

The HTTP Setup tab is displayed. See Table 63-13 on page 63-32 for a description of the fields on this tab.

Step 2 Select the check boxes to enable HTTP and SSL (HTTPS) server functionality on the router.

Note If SSL is disabled (or if the HTTP policy as a whole is unassigned), Security Manager cannot communicate with the device after deployment unless you change the transport protocol for this device to SSH. This setting can be found in Device Properties. See Managing Device Communication Settings and Certificates, page 9-4.

Tip We recommend that you disable HTTP when SSL is enabled. This is required to ensure only secure connections to the server.

Step 3 (Optional) Modify the default ports used by HTTP (80) and HTTPS (443).

Step 4 (Optional) In the Allow Connection From field, enter the name of the standard, numbered ACL object that specifies which addresses can use HTTP and HTTPS on this device, or click **Select** to select the ACL object from a list or to create a new one. Use this option to restrict access to these protocols. For more information about creating standard ACL objects, see Creating Standard Access Control List Objects, page 6-56

Note Make sure that the ACL you select permits the Security Manager server; otherwise, communication with the device is lost.

Step 5 (Optional) On the AAA tab, modify the default type of authentication to perform on users who attempt to access the device using HTTP or HTTPS. Options include AAA, Enable Password (default), Local Database, and TACACS.

If you select AAA, continue with Step 6; otherwise, continue with Step 8.

Note The TACACS option applies only to devices using an IOS software version prior to 12.3(8).

See Table 63-14 on page 63-33 for a description of the fields on the AAA tab.

Step 6 Select the authentication method to perform on users:

- If you want to use the default AAA login authentication methods defined in the device's AAA policy (see Defining AAA Services, page 63-4), do *not* select the Enable Device Login Authentication check box. Continue with Step 7.
- If you want to define a method list especially for this policy, do the following:

a. Select the **Enable Device Login Authentication** check box.

b. Under Prioritized Method List, enter the names of the AAA server groups to use for authentication, or click **Select** to select the AAA server groups from a list or to create new ones. Use the up and down arrows in the selector to define the order in which you want to apply these authentication methods.

Note Make sure that Security Manager users are defined on the AAA servers; otherwise communication with the device is lost.

Step 7 Select the authorization method to perform on users who use HTTP or HTTPS to begin an EXEC session:

- If you want to use the default AAA authorization methods defined in the device's AAA policy, do *not* select the Enable CLI/EXEC Operations Authorization check box. Continue with Step 8.
- If you want to define a method list especially for this policy, select the **Enable CLI/EXEC Operations Authorization** check box, then define the method list.

Note If you leave this option deselected, make sure that EXEC authorization is enabled in the router's AAA policy. Otherwise, you will be unable to connect to the device via HTTP or HTTPS (SSL). This applies to Security Manager as well as other applications, such as SDM. See Defining AAA Services, page 63-4.

Step 8 (Optional) Create command authorization definitions for specific privilege levels:

a. Click the **Add** button under the Command Authorization Override table. The Command Authorization Override dialog box is displayed. See Table 63-15 on page 63-34 for a description of the fields in this dialog box.

b. Configure the command authorization definition as required.

c. Click **OK**. The dialog box closes and the authorization method is displayed in the Command Authorization Override table.

d. Repeat a. through c. to create additional command authorization definitions.

HTTP Policy Page

Use the HTTP page to configure HTTP and HTTPS access on the router. You can configure HTTP policies on a Cisco IOS router from the following tabs on the HTTP policy page:

- HTTP Page—Setup Tab, page 63-31
- HTTP Page—AAA Tab, page 63-32

For more information, see HTTP and HTTPS on Cisco IOS Routers, page 63-28.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > HTTP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > HTTP** from the Policy Type selector. Right-click **HTTP** to create a policy, or select an existing policy from the Shared Policy selector.

HTTP Page—Setup Tab

Use the Setup tab of the HTTP page to enable HTTP and HTTP over Secure Socket Layer (HTTP over SSL or HTTPS) on the router. You can optionally limit access to these protocols to the addresses defined in an access control list.

Note As a general rule, Cisco IOS routers that have been discovered by Security Manager already have HTTPS enabled because Security Manager uses SSL as the default protocol for communicating with them. See Setting Up SSL on Cisco IOS Routers, page 2-4.

Navigation Path

Go to the HTTP Policy Page, page 63-31, then click the **Setup** tab.

Related Topics

- HTTP Page—AAA Tab, page 63-32
- HTTP and HTTPS on Cisco IOS Routers, page 63-28

Field Reference

Table 63-13 HTTP Page—Setup Tab

Element	Description
Enable HTTP	When selected, an HTTP server is enabled on the router. When deselected, HTTP is disabled on the router. This is the default for devices that were not discovered.
HTTP Port	The port number to use for HTTP. Valid values are 80 or any value from 1024 to 65535. The default is 80.
Enable SSL	When selected, a secure HTTP server (HTTP over SSL or HTTPS) is enabled on the router. When deselected, HTTPS is disabled. This is the default for devices that were not discovered. **Note** If SSL is disabled (or if the HTTP policy as a whole is unassigned), Security Manager cannot communicate with the device after deployment unless you change the transport protocol for this device to SSH. This setting can be found in Device Properties. **Note** We recommend that you disable HTTP when SSL is enabled. This is required to ensure only secure connections to the server.
SSL Port	The port number to use for HTTPS. Valid values are 443 or any value from 1025 to 65535. The default is 443.
Allow Connection From	The name of the standard numbered ACL that restricts use of HTTP and HTTPS on this device. Enter the name of an ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it. **Note** If you define an ACL, make sure that it includes the Security Manager server. Otherwise, Security Manager cannot communicate with this device using SSL.

HTTP Page—AAA Tab

Use the AAA tab of the HTTP page to define the authentication and authorization methods to perform on users who attempt to access the router using HTTP or HTTPS.

Navigation Path

Go to the HTTP Policy Page, page 63-31, then click the **AAA** tab.

Related Topics

- HTTP Page—Setup Tab, page 63-31
- HTTP and HTTPS on Cisco IOS Routers, page 63-28
- Filtering Tables, page 1-48

Field Reference

Table 63-14 HTTP Page—AAA Tab

Element	Description
Authenticate Using	The type of authentication to use: • AAA—Performs AAA login authentication. • Enable Password—Uses the enable password configured on the router. This is the default. • Local Database—Uses the local username database configured on the router. • TACACS—Uses the TACACS or XTACACS server configured on the router. Applies only to devices using an IOS software version prior to 12.3(8) or 12.3(8)T.
Login Authentication settings	
Enable Device Login Authentication	Applies only when AAA is selected as the authentication method. When selected, authentication is based on the methods defined in the Prioritized Method List field. When deselected, the default authentication list defined in the router's AAA policy is used. See AAA Page—Authentication Tab, page 63-6.
Prioritized Method List	Applies only when the Enable Device Login Authentication check box is selected. Defines a sequential list of methods to be queried when authenticating a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authenticate users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
EXEC Authorization settings	
Enable CLI/EXEC Operations Authorization	Applies only when AAA is selected as the authentication method. When selected, EXEC authorization is based on the methods defined in the Prioritized Method List field. This type of authorization determines whether the user is permitted to open an EXEC (CLI) session. When deselected, the default EXEC authorization list defined in the router's AAA policy is used. See AAA Page—Authorization Tab, page 63-8. **Note** If you leave this option deselected, make sure that EXEC authorization is enabled in the router's AAA policy. Otherwise, you will be unable to connect to the device via HTTP or HTTPS (SSL). This applies to Security Manager as well as other applications, such as SDM and the device's web interface.

Table 63-14 HTTP Page—AAA Tab (continued)

Element	Description
Prioritized Method List	Applies only when the Enable CLI/EXEC Operations Authorization check box is selected. Defines a sequential list of methods to be queried when authorizing a user to open an EXEC (CLI) session. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Command Authorization settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Authorization Override Dialog Box, page 63-34. From here you can configure a command authorization definition.
Edit button	Opens the Command Authorization Override Dialog Box, page 63-34. From here you can edit the command authorization definition.
Delete button	Deletes the selected command authorization definitions from the table.

Command Authorization Override Dialog Box

Use the Command Authorization Override dialog box to define which methods to use when authorizing the EXEC commands that are associated with a given privilege. This enables you to authorize all commands associated with a specific privilege level, from 0 to 15.

Navigation Path

From the HTTP Page—AAA Tab, page 63-32, click the **Add** button beneath the Command Authorization Override table.

Related Topics

- HTTP Policy Page, page 63-31
- AAA Policy Page, page 63-6

Field Reference

Table 63-15 Command Authorization Dialog Box

Element	Description
Privilege Level	The privilege level for which you want to define a command accounting list. Valid values range from 0 to 15.

Table 63-15 Command Authorization Dialog Box (continued)

Element	Description
Prioritized Method List	Defines a sequential list of methods to be used when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. Supported methods include TACACS+, Local, and None. **Note** If you select None as a method, it must appear as the last method in the list.

Line Access on Cisco IOS Routers

Security Manager enables you to configure command line access (also called EXEC access) to a router using the following methods:

- Console port—Physical connection via a standard RS232 cable for local access. For more information, see:
  - Defining Console Port Setup Parameters, page 63-35
  - Defining Console Port AAA Settings, page 63-37
- VTY lines—Virtual terminal lines for remote access, typically using protocols such as Telnet, SSH, or rlogin. For more information, see:
  - Defining VTY Line Setup Parameters, page 63-38
  - Defining VTY Line AAA Settings, page 63-40

After you configure and deploy these policies, you can use these lines to communicate with individual devices directly when you want to configure or diagnose them using the CLI.

Defining Console Port Setup Parameters

The console port on a router is generally used for local system access by an administrator with physical access to the device. By default, the console port is set up as follows:

- All permitted users have privileged access to the router, including all configuration commands (privilege level 15).
- The line is disconnected after 10 minutes without user input.
- Incoming connections are not permitted.
- Outgoing connections support Telnet only.

In addition to modifying any of the default settings, you can optionally define the following settings:

- The password for accessing the console.
- Whether to disable all EXEC sessions on the console.

- Incoming and outgoing ACLs that restrict the connections that are permitted on the console.
- Whether VRF connections are permitted on the console.

Related Topics

- Line Access on Cisco IOS Routers, page 63-35

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > Console** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > Console** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The Console Setup tab is displayed. See Table 63-16 on page 63-42 for a description of the fields on this tab.

Step 2 (Optional) Enter the password for accessing the console port, then enter it again in the Confirm field.

Step 3 (Optional) Modify the default (15) granted to users of the console port. See Console Page—Authorization Tab, page 63-45.

Step 4 (Optional) Select the **Disable all the EXEC sessions to the router via this line** check box to prevent any incoming connections via the console.

Note Selecting this option blocks all access to the device via the console port.

Step 5 (Optional) Modify the default timeout after which the line is disconnected if no user input is detected.

Note Setting this value to 0 disables the timeout. Disabling the timeout could compromise the security of your network.

Step 6 (Optional) Specify which protocols can be used for outbound connections on the console port:

- All—All supported protocols are permitted.
- None—No protocols are permitted.
- Protocol—Enables one or more of the following protocols: SSH, Telnet, and rlogin.

Note You must configure AAA authentication on devices where the console port permits the SSH and rlogin protocols. See Defining Console Port AAA Settings, page 63-37.

Step 7 (Optional) Enter the names of ACLs that restrict incoming and outgoing connections between the device and the addresses in these lists, or click **Select** to select the ACL object or to create a new one. At the top of the selector, in the Type field, select the ACL type as either Standard or Extended.

Step 8 (Optional) Click the **AAA** tab to define authentication, authorization, and accounting settings for the console port. See Defining Console Port AAA Settings, page 63-37.

Defining Console Port AAA Settings

By default, authentication, authorization, and accounting are not performed on the console port. When you configure one or more of these access control options, you can either make use of the default method lists defined in the device's AAA policy or define a custom method list containing one or more AAA methods.

Related Topics

- Defining Console Port Setup Parameters, page 63-35
- Line Access on Cisco IOS Routers, page 63-35

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > Console** from the Policy selector, then click the **Authentication** tab in the work area.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > Console** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Authentication** tab.

The Console Authentication tab is displayed.

Step 2 (Optional) Select the authentication method to perform on users who attempt to access the console line.

See Table 63-17 on page 63-45 for a description of the fields on the Authentication tab.

Note If you select local authentication, preview the full configuration before deployment to make sure that the **aaa new-model** command is not configured by another policy (for example, by configuring a method list in the AAA policy) or is already configured on the device itself.

Step 3 (Optional) On the Authorization tab, select the authorization method to perform on users who access the console line and begin an EXEC session.

See Table 63-18 on page 63-46 for a description of the fields on the Authorization tab.

Note RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization.

Step 4 (Optional) Create command authorization definitions for specific privilege levels:

a. Click the **Add** button under the Commands Authorization table. The Command Authorization dialog box is displayed. See Table 63-26 on page 63-61 for details.

b. Configure the command authorization definition as required.

c. Click **OK**. The dialog box closes and the authorization method is displayed in the Commands Authorization table.

d. Repeat a. through c. to create additional command authorization definitions.

Step 5 (Optional) On the Accounting tab, select the EXEC and connection accounting methods to perform on users who access the console line.

See Table 63-19 on page 63-47 for a description of the fields on this tab.

Step 6 (Optional) Create command accounting definitions for specific privilege levels:

a. Click the **Add** button under the Commands Accounting table. The Command Accounting Dialog Box—Line Access, page 63-61 is displayed.

b. Configure the command accounting definition as required.

c. Click **OK**. The dialog box closes and the accounting method is displayed in the Commands Accounting table.

d. Repeat a. through c. to create additional command accounting definitions.

Defining VTY Line Setup Parameters

All Cisco IOS routers are configured by default with five VTY lines (labeled 0-4) that have the following settings:

- All permitted users have privileged access to the router, including all configuration commands (privilege level 15).
- VTY lines are disconnected after 10 minutes without user input.
- Incoming connections are not permitted.
- Outgoing connections support Telnet only.

You can use Security Manager to modify the default settings on these five VTY lines or to configure additional lines (up to a maximum of 16). In addition, you can optionally configure the following settings on each line:

- The password for accessing the line.
- Whether to disable all EXEC sessions on the line.
- Incoming and outgoing ACLs that restrict the connections that are permitted on the line.
- Whether VRF connections are permitted on the line.

Defining Groups of VTY Lines

You can configure multiple VTY lines as a contiguous group, which enables you to define identical settings for all the lines in the group with one procedure. All the lines within the group must fall within one of two ranges, 0-4 or 6-15. The group cannot overlap these two ranges.

The rules for configuring VTY line 5 are as follows. Line 5 can be part of the same definition as lines 0-4 only when there are no lines configured above line 5. If there are lines configured above line 5, you cannot include line 5 in the definition for lines 0-4, even if their configurations are the same. Line 5 *can* be included in the definition of the lines above line 5 if their configurations are the same.

For example, if lines 0-5 all share one configuration and lines 6-9 have a different configuration, you need to create three definitions—one definition for lines 0-4, a second definition for line 5, and a third definition for lines 6-9.

Note When you configure VTY lines, bear in mind that users are assigned a line at random when they connect to the device.

Note You can create only one definition per VTY line. An error is displayed if you create a VTY line definition that overlaps an existing definition.

Note If you use Security Manager to configure the default VTY lines (0-4), your definition overrides the default settings on the device. If you later delete this definition from Security Manager, the input protocol settings are retained and the other default settings are restored. This ensures that you always have VTY lines available for remote access to the device.

Note You can use the CLI or FlexConfigs to configure additional VTY lines on devices that support more than 16 lines.

Related Topics

- Line Access on Cisco IOS Routers, page 63-35

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > VTY** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > VTY** from the Policy Type selector. Select an existing policy or create a new one.

The VTY page is displayed. See Table 63-20 on page 63-50 for a description of the fields on this page.

Step 2 Click the **Add** button beneath the Lines table, or select a line definition and then click the **Edit** button. The Setup tab of the VTY Lines dialog box is displayed. See Table 63-21 on page 63-51 for a description of the fields on this tab.

Step 3 Enter the relative line number of the VTY line. If you are configuring a group of VTY lines, enter the first and last numbers of the group in the fields provided.

Step 4 (Optional) Enter the password for accessing the console line, then enter it again in the Confirm field.

Step 5 (Optional) Modify the default Privilege (15) granted to users of this VTY line (or group of lines).

Step 6 (Optional) Select the **Disable all the EXEC sessions to the router via this line** check box to prevent any incoming connections over this VTY line (or group of lines).

Step 7 (Optional) Modify the default timeout after which the line is disconnected if no user input is detected.

Note Setting this value to 0 to disables the timeout. Disabling the timeout could cause abandoned sessions to block available VTY lines. It can also compromise the security of your network.

Step 8 (Optional) Specify which protocols can be used for inbound and outbound connections on this VTY line (or group of lines):

- All—All supported protocols are permitted.
- None—No protocols are permitted.
- Protocol—Enables one or more of the following protocols: SSH, Telnet, and rlogin.

Caution Setting the inbound connections setting to None might prevent Security Manager from connecting to the device after deployment.

Note You must configure AAA authentication when the VTY line permits the SSH and rlogin protocols. See Defining VTY Line AAA Settings, page 63-40.

Step 9 (Optional) Enter the names of ACLs that restrict incoming and outgoing connections between the device and the addresses in these lists, or click **Select** to select an ACL object from a list or to create a new one. You can choose from standard or extended ACLs.

Tip Defining an inbound ACL is a good way to reserve a VTY line for administrative access only.

Step 10 (Optional) Click the **AAA** tab to define authentication, authorization, and accounting settings for this VTY line (or group of lines). See Defining VTY Line AAA Settings, page 63-40.

Step 11 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the Lines table.

Note To remove a VTY line definition, select it, then click **Delete**. If you delete a VTY line from an IOS device, any subsequent lines are also deleted. For example, if the device contains lines 0-9 and you delete line 5, lines 6-9 are deleted as well. If you delete the definition for lines 0-4 from Security Manager, the router retains the inbound protocol definition and restores the other default settings for these lines on the device. This ensures that five VTY lines are always available.

Defining VTY Line AAA Settings

By default, authentication, authorization, and accounting are not performed on VTY lines. When you configure one or more of these access control options, you can either make use of the default method lists defined in the device's AAA policy or define a custom method list containing one or more AAA methods.

Before You Begin

- Define the basic parameters of the VTY line or group of VTY lines. See Defining VTY Line Setup Parameters, page 63-38.

Related Topics

- Defining VTY Line Setup Parameters, page 63-38
- Line Access on Cisco IOS Routers, page 63-35

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > VTY** from the Policy selector.

- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > VTY** from the Policy Type selector. Select an existing policy or create a new one.

The VTY page is displayed. See Table 63-20 on page 63-50 for a description of the fields on this page.

Step 2 Select a VTY line definition in the Lines tables, click the **Edit** button to display the VTY Line dialog box, then click the **Authentication** tab.

Step 3 (Optional) Select the authentication method to perform on users who attempt to access the VTY line.

See Table 63-23 on page 63-55 for a description of the fields on this tab.

Note If you select local authentication, preview the full configuration before deployment to make sure that the **aaa new-model** command is not configured by another policy (for example, by configuring a method list in the AAA policy) or is already configured on the device itself.

Step 4 (Optional) On the Authorization tab, select the authorization method to perform on users who access the VTY line and begin an EXEC session.

See Table 63-24 on page 63-56 for a description of the fields on the Authorization tab.

Note RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization.

Step 5 (Optional) Create command authorization definitions for specific privilege levels:

a. Click the **Add** button under the Commands Authorization table. The Command Authorization Dialog Box—Line Access, page 63-60 is displayed.

b. Configure the command authorization definition as required.

c. Click **OK**. The dialog box closes and the authorization method is displayed in the Commands Authorization table.

d. Repeat a. through c. to create additional command authorization definitions.

Step 6 (Optional) On the Accounting tab, select the EXEC and connection accounting methods to perform on users who attempt to access the VTY line.

See Table 63-25 on page 63-58 for a description of the fields on the Accounting tab.

Step 7 (Optional) Create command accounting definitions for specific privilege levels:

a. Click the **Add** button under the Commands Accounting table. The Command Accounting Dialog Box—Line Access, page 63-61 is displayed.

b. Configure the command accounting definition as required.

c. Click **OK**. The dialog box closes and the accounting method is displayed in the Commands Accounting table.

d. Repeat a. through c. to create additional command accounting definitions.

Console Policy Page

Use the Console page to configure access to the router over the console port. You can configure console policies on a Cisco IOS router from the following tabs on the Console policy page:

- Console Page—Setup Tab, page 63-42
- Console Page—Authentication Tab, page 63-44
- Console Page—Authorization Tab, page 63-45
- Console Page—Accounting Tab, page 63-47

For more information, see Line Access on Cisco IOS Routers, page 63-35.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > Console** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > Console** from the Policy Type selector. Right-click **Console** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- VTY Policy Page, page 63-50

Console Page—Setup Tab

Use the Setup tab of the Console page to define the basic parameters of the console port. This includes the password for accessing the port, the privilege level assigned to users, the protocols that are permitted, and the ACLs that limit access.

Navigation Path

Go to the Console Policy Page, page 63-42, then click the **Setup** tab.

Related Topics

- Console Page—Authentication Tab, page 63-44
- Console Page—Authorization Tab, page 63-45
- Console Page—Accounting Tab, page 63-47
- VTY Line Dialog Box—Setup Tab, page 63-52

Field Reference

Table 63-16 Console Page—Setup Tab

Element	Description
Password	The password for accessing the console port. The password is case sensitive and can contain up to 80 alphanumeric characters. The first character cannot be a number. Spaces are not allowed. Enter the password again in the Confirm field.

Table 63-16 Console Page—Setup Tab (continued)

Element	Description
Privilege Level	The privilege level assigned to users connected to the console port. Valid values range from 0 to 15: • 0—Grants access to these commands only: **disable**, **enable**, **exit**, **help**, and **logout**. • 1—Enables nonprivileged access to the router (normal EXEC-mode use privileges). • 15—Enables privileged access to the router (traditional enable privileges). **Note** Levels 2-14 are not normally used in a default configuration, but custom configurations can be created by moving commands that are normally at level 15 to a lower level and commands that are normally at level 1 to a higher level. You can configure the privilege levels of commands using the CLI or by defining a FlexConfig. **Note** If you do not define a value, level 1 is assigned by default. This value does not appear in the device configuration.
Disable all the EXEC sessions to the router via this line	When selected, disables EXEC sessions over this line. Select this option when you want to allow only an outgoing connection on the console. This option is useful for keeping the console port free from unsolicited incoming data that can tie up the line. When deselected, EXEC sessions are enabled on the console port. This is the default. **Note** Selecting this option blocks all access to the device via the console port.
Exec Timeout	The amount of time (in seconds) that the EXEC command interpreter waits to detect user input on the console port. If no input is detected, the line is disconnected. Valid values range from 0 to 2147483. The default is 600 (10 minutes). Setting the value to 0 disables the timeout. **Note** Although the timeout is defined in seconds, it appears in the CLI in the format [mm ss].

Table 63-16 Console Page—Setup Tab (continued)

Element	Description
Output Protocols	The protocols that you can use for outgoing connections on the console port: • All—All supported protocols are permitted. Supported protocols include LAT, MOP, NASI, PAD, rlogin, SSH, Telnet, and V.120. • None—No protocols are permitted. This makes the port unusable by outgoing connections. • Protocol—Enables one or more of the following protocols: – SSH—Secure Shell protocol. – Telnet—Standard TCP/IP terminal emulation protocol. – rlogin—UNIX rlogin protocol. **Note** SSH and rlogin require that you configure AAA authentication. See Console Page—Authentication Tab, page 63-44. **Note** Not all IOS Software Versions support rlogin as an output protocol.
Inbound Access List	The name of the ACL object that restricts incoming connections on the console port. Enter the name of the ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Permit VRF Interface Connections	Applies only when an inbound ACL is defined on the console port. When selected, accepts incoming connections from interfaces that belong to a VRF. When deselected, rejects incoming connections from interfaces that belong to a VRF.
Outbound Access List	The name of the ACL object that restricts outgoing connections on the console port. Enter the name of an ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.

Console Page—Authentication Tab

Use the Authentication tab of the Console page to define the AAA authentication methods to perform on users who attempt to access the console port.

Navigation Path

Go to the Console Policy Page, page 63-42, then click the **Authentication** tab.

Related Topics

- Console Page—Setup Tab, page 63-42
- Console Page—Authorization Tab, page 63-45
- Console Page—Accounting Tab, page 63-47
- VTY Line Dialog Box—Authentication Tab, page 63-55

Field Reference

Table 63-17 Console Page—Authentication Tab

Element	Description
Authenticate Using	Authentication settings for the console port: • None—Authentication is not performed. This is the default. • Local Database—Uses the local username database for authentication. • AAA Policy Default List—Uses the default authentication method list that is defined in the device's AAA policy. See AAA Page—Authentication Tab, page 63-6. • Custom Method List—Uses the authentication methods specified in the Authentication Method List field. **Note** If you select local authentication, preview the full configuration before deployment to make sure that the **aaa new-model** command is not configured by another policy (for example, by configuring a method list in the AAA policy) or is already configured on the device itself.
Prioritized Method List	Applies only when Custom Method List is selected as the authentication method. Defines a sequential list of methods to be queried when authenticating a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authenticate users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.

Console Page—Authorization Tab

Use the Authorization tab of the Console page to define the EXEC and command authorization methods to perform on users who access the console port.

Note You must enable AAA services on the router to use this feature; otherwise, deployment will fail. See Defining AAA Services, page 63-4.

Navigation Path

Go to the Console Policy Page, page 63-42, then click the **Authorization** tab.

Related Topics

- Console Page—Setup Tab, page 63-42

- Console Page—Authentication Tab, page 63-44
- Console Page—Accounting Tab, page 63-47
- VTY Line Dialog Box—Authorization Tab, page 63-56
- Filtering Tables, page 1-48

Field Reference

Table 63-18 Console Page—Authorization Tab

Element	Description
EXEC Authorization settings	
Authorize EXEC Operations Using	The authorization method that determines whether a user is allowed to run an EXEC session: • None—Authorization is not performed. This is the default. • AAA Policy Default List—Uses the default authorization method list that is defined in the device's AAA policy. See AAA Page—Authorization Tab, page 63-8. • Custom Method List—Uses the authorization methods specified in the EXEC Method List field.
Prioritized Method List	Applies only when Custom Method List is selected as the EXEC method. Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list. **Note** RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization.
Command Authorization settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Authorization Dialog Box—Line Access, page 63-60. From here you can configure a command authorization definition.

Table 63-18 Console Page—Authorization Tab (continued)

Element	Description
Edit button	Opens the Command Authorization Dialog Box—Line Access, page 63-60. From here you can edit the command authorization definition.
Delete button	Deletes the selected command authorization definitions from the table.

Console Page—Accounting Tab

Use the Accounting tab of the Console page to define the EXEC, connection, and command accounting methods to perform on users who access the console port.

Note You must enable AAA services on the router to use this feature; otherwise, deployment will fail. See Defining AAA Services, page 63-4.

Navigation Path

Go to the Console Policy Page, page 63-42, then click the **Accounting** tab.

Related Topics

- Console Page—Setup Tab, page 63-42
- Console Page—Authentication Tab, page 63-44
- Console Page—Authorization Tab, page 63-45
- VTY Line Dialog Box—Accounting Tab, page 63-57
- Filtering Tables, page 1-48

Field Reference

Table 63-19 Console Page—Accounting Tab

Element	Description
EXEC Accounting settings	
Perform EXEC Accounting Using	The accounting method to use for recording basic information about user EXEC sessions: • None—Accounting is not performed. This is the default. • AAA Policy Default List—Uses the default EXEC accounting method list that is defined in the device's AAA policy. See AAA Page—Accounting Tab, page 63-10. • Custom Method List—Uses the accounting methods specified in the EXEC Method List field. EXEC accounting records basic details about EXEC sessions, such as the username, date, start and stop times, and the access server IP address.

Table 63-19 Console Page—Accounting Tab (continued)

Element	Description
Generate Accounting Records for	Applies only when Custom Method List is selected as the EXEC method. Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. This is the default. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.
Prioritized Method List	Applies only when Custom Method List is selected as the EXEC method. Defines a sequential list of methods to be queried when creating accounting methods for a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Enable Broadcast to Multiple Servers	Applies only when Method List is selected as the EXEC method. When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.
Connection Accounting settings	
Perform Connection Accounting Using	The accounting method to use for recording information about outbound connections made over the console line: • None—Accounting is not performed. This is the default. • AAA Policy Default List—Uses the default connection accounting method list that is defined in the device's AAA policy. See AAA Page—Accounting Tab, page 63-10. • Custom Method List—Uses the accounting methods specified in the Connection Method List field. Connection accounting records details about outgoing connections over the line, such as Telnet and rlogin connections.

Table 63-19 Console Page—Accounting Tab (continued)

Element	Description
Generate Accounting Records for	Applies only when Custom Method List is selected as the connection method. Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. This is the default. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.
Prioritized Method List	Applies only when Custom Method List is selected as the connection method. Defines a sequential list of methods to be queried when creating accounting methods for a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Enable Broadcast to Multiple Servers	Applies only when Custom Method List is selected as the connection method. When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.
Command Accounting settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Generate Accounting Records for	The points in the process where the device sends an accounting notice to the accounting server.
Enable Broadcast	Whether accounting records are broadcast to multiple servers simultaneously.
Prioritized Method List	The method list to use when authorizing users with this privilege level.

Table 63-19 Console Page—Accounting Tab (continued)

Element	Description
Add button	Opens the Command Accounting Dialog Box—Line Access, page 63-61. From here you can configure a command accounting definition.
Edit button	Opens the Command Accounting Dialog Box—Line Access, page 63-61. From here you can edit the command accounting definition.
Delete button	Deletes the selected command accounting definitions from the table.

VTY Policy Page

Use the VTY page to configure up to 16 VTY lines for remote access to the router. In addition to configuring individual lines, you can configure a group of lines that share the same definition.

For more information, see Line Access on Cisco IOS Routers, page 63-35.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Line Access > VTY** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Line Access > VTY** from the Policy Type selector. Right-click **VTY** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Console Policy Page, page 63-42
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-20 VTY Lines Page

Element	Description
Line	The relative line number of the VTY line. This field may also contain multiple VTY lines configured as a contiguous group.
Line/Line Group Parameters	
Input Protocols	The protocols that you can use for incoming connections on the VTY line.
Output Protocols	The protocols that you can use for outgoing connections on the VTY line.
Privilege Level	The privilege level assigned to users.
Exec Timeout	The amount of time the EXEC command interpreter waits until user input is detected.
Inbound ACL	The ACL used to limit inbound traffic.

Table 63-20 VTY Lines Page (continued)

Element	Description
Outbound ACL	The ACL used to limit outbound traffic.
Authentication	The type of AAA authentication used.
Authorization	The types of AAA authorization used.
Accounting	The types of AAA accounting used.
VTY Line Page Buttons	
Add button	Opens the VTY Line Dialog Box, page 63-51. From here you can define a VTY line or line group.
Edit button	Opens the VTY Line Dialog Box, page 63-51. From here you can edit the VTY line or line group.
Delete button	Deletes the selected VTY lines from the table. If you delete a VTY line from an IOS device, any subsequent lines are also deleted. For example, if the device contains lines 0-9 and you delete line 5, lines 6-9 are deleted as well. **Note** If you delete any of the default VTY lines (0-4) on the device, the input protocol settings are retained and the other default settings are restored. This helps prevent you from cutting off remote access to the device.

VTY Line Dialog Box

Use the VTY Line dialog box to configure one or more VTY lines (up to 16) that enable remote users to access the router. When you configure a VTY line, you can define the type of authentication and authorization to perform on users who access the lines.

Navigation Path

Go to the VTY Policy Page, page 63-50, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Line Access on Cisco IOS Routers, page 63-35
- Console Policy Page, page 63-42

Field Reference

Table 63-21 VTY Line Dialog Box

Element	Description
Setup tab	Defines the basic configuration of the VTY line or line group. See VTY Line Dialog Box—Setup Tab, page 63-52.
Authentication tab	Defines the type of AAA authentication to perform on users who access the VTY line. See VTY Line Dialog Box—Authentication Tab, page 63-55.

Table 63-21 VTY Line Dialog Box (continued)

Element	Description
Authorization tab	Defines the types of AAA authorization to perform on users who access the VTY line. See VTY Line Dialog Box—Authorization Tab, page 63-56.
Accounting tab	Defines the types of AAA accounting to perform on users who access the VTY line. See VTY Line Dialog Box—Accounting Tab, page 63-57.

VTY Line Dialog Box—Setup Tab

Use the Setup tab of the VTY Line dialog box to define the basic parameters of the VTY line. This includes the password for accessing the line, the privilege level assigned to users, the protocols that are permitted on the line, and the ACLs that limit access.

Navigation Path

Go to the VTY Line Dialog Box, page 63-51, then click the **Setup** tab.

Related Topics

- Defining VTY Line Setup Parameters, page 63-38
- VTY Line Dialog Box—Authentication Tab, page 63-55
- VTY Line Dialog Box—Authorization Tab, page 63-56
- VTY Line Dialog Box—Accounting Tab, page 63-57
- Console Page—Setup Tab, page 63-42

Field Reference

Table 63-22 VTY Line Dialog Box—Setup Tab

Element	Description
Starting VTY Line Number	The relative line number of the VTY line. If you are configuring a group of VTY lines, enter the number of the first line in the group. Valid values range from 0 to 15. **Note** Although different routers support a different number of VTY lines (from four to several thousand), Security Manager supports a maximum of 16 lines per device. You cannot configure the same line number more than once.
Ending VTY Line Number	Applies only when configuring a group of lines. The relative line number of the last VTY line in the group. **Note** When you configure a group of lines, all the lines in the group must fall within one of two ranges, 0-4 or 6-15.

Table 63-22 VTY Line Dialog Box—Setup Tab (continued)

Element	Description
Password	The password for accessing this VTY line. The password is case sensitive and can contain up to 80 alphanumeric characters. The first character cannot be a number. Spaces are not allowed. Enter the password again in the Confirm field.
Privilege Level	The privilege level assigned to users on this VTY line. Valid values range from 0 to 15: • 0—Grants access to these commands only: **disable**, **enable**, **exit**, **help**, and **logout**. • 1—Enables nonprivileged access to the router (normal EXEC-mode use privileges). • 15—Enables privileged access to the router (traditional enable privileges). **Note** Levels 2-14 are not normally used in a default configuration, but custom configurations can be created by moving commands that are normally at level 15 to a lower level and commands that are normally at level 1 to a higher level. You can configure the privilege levels of commands using the CLI or by defining a FlexConfig. **Note** If you do not define a value, level 1 is assigned by default. This value does not appear in the device configuration.
Disable all the EXEC sessions to the router via this line	When selected, EXEC sessions are disabled over this line. Select this option when you want to allow only an outgoing connection on this line. This option is useful for keeping a particular line free from unsolicited incoming data that can tie up the line. When deselected, EXEC sessions are enabled over this line. This is the default.
Exec Timeout	The amount of time (in seconds) that the EXEC command interpreter waits to detect user input on the line. If no input is detected, the line is disconnected. Valid values range from 0 to 2147483. The default is 600 (10 minutes). Setting the value to 0 disables the timeout. **Note** Although the timeout is defined in seconds, it appears in the CLI in the format [mm ss].

Table 63-22 VTY Line Dialog Box—Setup Tab (continued)

Element	Description
Input Protocols	The protocols that you can use for incoming connections on this line: • All—All supported protocols are permitted. Supported protocols include LAT, MOP, NASI, PAD, rlogin, SSH, Telnet, and V.120. • None—No protocols are permitted. This makes the port unusable by incoming SSH, Telnet, and rlogin connections. **Note** Setting the input protocols setting to None might prevent Security Manager from connecting to the device after deployment. The device can still be managed using SSL, if SSL is enabled in the HTTP policy. See HTTP Page—Setup Tab, page 63-31. • Protocol—Enables one or more of the following protocols: – SSH—Secure Shell protocol. – Telnet—Standard TCP/IP terminal emulation protocol. – rlogin—UNIX rlogin protocol. **Note** SSH and rlogin require that you configure AAA authentication. See VTY Line Dialog Box—Authentication Tab, page 63-55. **Note** Not all IOS Software Versions support rlogin as an input protocol.
Output Protocols	The protocols that you can use for outgoing connections on this line: • All—All supported protocols are permitted. Supported protocols include LAT, MOP, NASI, PAD, rlogin, SSH, Telnet, and V.120. • None—No protocols are permitted. This makes the port unusable by outgoing connections. • Protocol—Enables one or more of the following protocols: – SSH—Secure Shell protocol. – Telnet—Standard TCP/IP terminal emulation protocol. – rlogin—UNIX rlogin protocol. **Note** SSH and rlogin require that you configure AAA authentication. See VTY Line Dialog Box—Authentication Tab, page 63-55. **Note** Not all IOS Software Versions support rlogin as an output protocol.
Inbound Access List	The name of the ACL object that restricts incoming connections on this line. Enter the name of the ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.
Permit VRF Interface Connections	Applies only when an inbound ACL is defined on this line. When selected, accepts incoming connections from interfaces that belong to a VRF. When deselected, rejects incoming connections from interfaces that belong to a VRF.

Table 63-22 VTY Line Dialog Box—Setup Tab (continued)

Element	Description
Outbound Access List	The name of the ACL object that restricts outgoing connections on this line. Enter the name of the ACL object, or click **Select** to select it. If the object that you want is not listed, click the **Create** button to create it.

VTY Line Dialog Box—Authentication Tab

Use the Authentication tab of the VTY Line dialog box to define the authentication methods to perform on users who attempt to access the selected VTY line or group of lines.

Navigation Path

Go to the VTY Line Dialog Box, page 63-51, then click the **Authentication** tab.

Related Topics

- Defining VTY Line AAA Settings, page 63-40
- VTY Line Dialog Box—Setup Tab, page 63-52
- VTY Line Dialog Box—Authorization Tab, page 63-56
- VTY Line Dialog Box—Accounting Tab, page 63-57
- Console Page—Authentication Tab, page 63-44

Field Reference

Table 63-23 VTY Line Dialog Box—Authentication Tab

Element	Description
Authenticate Using	Authentication settings for the VTY line: • None—Authentication is not performed. This is the default. • Local Database—Uses the local username database for authentication. • AAA Policy Default List—Uses the default authentication method list that is defined in the device's AAA policy. See AAA Page—Authentication Tab, page 63-6. • Custom Method List—Uses the authentication methods specified in the Prioritized Method List field. **Note** If you select local authentication, preview the full configuration before deployment to make sure that the **aaa new-model** command is not configured by another policy (for example, by configuring a method list in the AAA policy) or is already configured on the device itself.

Table 63-23 VTY Line Dialog Box—Authentication Tab (continued)

Element	Description
Prioritized Method List	Applies only when Custom Method List is selected as the authentication method. Defines a sequential list of methods to be queried when authenticating a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authenticate users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.

VTY Line Dialog Box—Authorization Tab

Use the Authorization tab of the VTY Line dialog box to define the EXEC and command authorization methods to perform on users who access the selected VTY line or group of lines.

Note You must enable AAA services on the router to use this feature; otherwise, deployment will fail. See Defining AAA Services, page 63-4.

Navigation Path

Go to the VTY Line Dialog Box, page 63-51, then click the **Authorization** tab.

Related Topics

- Defining VTY Line AAA Settings, page 63-40
- VTY Line Dialog Box—Setup Tab, page 63-52
- VTY Line Dialog Box—Authentication Tab, page 63-55
- VTY Line Dialog Box—Accounting Tab, page 63-57
- Console Page—Authentication Tab, page 63-44
- Filtering Tables, page 1-48

Field Reference

Table 63-24 VTY Line Dialog Box—Authorization Tab

Element	Description
EXEC Authorization settings	

Table 63-24 VTY Line Dialog Box—Authorization Tab (continued)

Element	Description
Authorize EXEC Operations Using	The authorization method that determines whether a user is allowed to run an EXEC session: • None—Authorization is not performed. This is the default. • AAA Policy Default List—Uses the default authorization method list that is defined in the device's AAA policy. See AAA Page—Authorization Tab, page 63-8. • Custom Method List—Uses the authorization methods specified in the Prioritized Method List field.
Prioritized Method List	Applies only when Custom Method List is selected as the EXEC method. Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list. **Note** RADIUS uses the same server for authentication and authorization. Therefore, if you use define a RADIUS method list for authentication, you must define the same method list for authorization.
Command Authorization settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Authorization Dialog Box—Line Access, page 63-60. From here you can configure a command authorization definition.
Edit button	Opens the Command Authorization Dialog Box—Line Access, page 63-60. From here you can edit the command authorization definition.
Delete button	Deletes the selected command authorization definitions from the table.

VTY Line Dialog Box—Accounting Tab

Use the Accounting tab of the VTY Line dialog box to define the EXEC, connection, and command accounting methods to perform on users who access the selected VTY line or group of lines.

Note You must enable AAA services on the router to use this feature; otherwise, deployment will fail. See Defining AAA Services, page 63-4.

Navigation Path

Go to the VTY Line Dialog Box, page 63-51, then click the **Accounting** tab.

Related Topics

- Defining VTY Line AAA Settings, page 63-40
- VTY Line Dialog Box—Setup Tab, page 63-52
- VTY Line Dialog Box—Authentication Tab, page 63-55
- Console Page—Accounting Tab, page 63-47
- Filtering Tables, page 1-48

Field Reference

Table 63-25 VTY Line Dialog Box—Accounting Tab

Element	Description
EXEC Accounting settings	
Perform EXEC Accounting Using	The accounting method to use for recording basic information about user EXEC sessions: • None—Accounting is not performed. This is the default. • AAA Policy Default List—Uses the default EXEC accounting method list that is defined in the device's AAA policy. See AAA Page—Accounting Tab, page 63-10. • Custom Method List—Uses the accounting methods specified in the Prioritized Method List field. EXEC accounting records basic details about EXEC sessions, such as the username, date, start and stop times, and the access server IP address.
Generate Accounting Records for	Applies only when Custom Method List is selected as the EXEC method. Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. This is the default. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.

Table 63-25 VTY Line Dialog Box—Accounting Tab (continued)

Element	Description
Prioritized Method List	Applies only when Custom Method List is selected as the EXEC method. Defines a sequential list of methods to be queried when creating accounting methods for a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Enable Broadcast to Multiple Servers	Applies only when Method List is selected as the EXEC method. When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.
Connection Accounting settings	
Perform Connection Accounting Using	The accounting method to use for recording information about outbound connections made over the VTY line: • None—Accounting is not performed. This is the default. • AAA Policy Default List—Uses the default connection accounting method list that is defined in the device's AAA policy. See AAA Page—Accounting Tab, page 63-10. • Custom Method List—Uses the accounting methods specified in the Prioritized Method List field. Connection accounting records details about outgoing connections over the line, such as Telnet and rlogin connections.
Generate Accounting Records for	Applies only when Custom Method List is selected as the connection method. Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. This is the default. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.

Table 63-25 VTY Line Dialog Box—Accounting Tab (continued)

Element	Description
Prioritized Method List	Applies only when Custom Method List is selected as the connection method. Defines a sequential list of methods to be queried when creating accounting methods for a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Enable Broadcast to Multiple Servers	Applies only when Custom Method List is selected as the connection method. When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.
Command Accounting settings	
Privilege Level	The privilege level to which the command authorization definition applies.
Generate Accounting Records for	The points in the process where the device sends an accounting notice to the accounting server.
Enable Broadcast	Whether accounting records are broadcast to multiple servers simultaneously.
Prioritized Method List	The method list to use when authorizing users with this privilege level.
Add button	Opens the Command Accounting Dialog Box—Line Access, page 63-61. From here you can configure a command accounting definition.
Edit button	Opens the Command Accounting Dialog Box—Line Access, page 63-61. From here you can edit the command accounting definition.
Delete button	Deletes the selected command accounting definitions from the table.

Command Authorization Dialog Box—Line Access

Use the Command Authorization dialog box to define which methods to use when authorizing the EXEC commands that are associated with a given privilege. This enables you to authorize all commands associated with a specific privilege level, from 0 to 15.

Navigation Path

From the Console Page—Authorization Tab, page 63-45 or the VTY Line Dialog Box—Authorization Tab, page 63-56, click the **Add** button beneath the Command Authorization table.

Related Topics

- Console Policy Page, page 63-42
- VTY Policy Page, page 63-50

Field Reference

Table 63-26 Command Authorization Dialog Box—Line Access

Element	Description
Privilege Level	The privilege level for which you want to define a command authorization list. Valid values range from 0 to 15. **Note** If you do not define a value, level 1 is assigned by default. This value does not appear in the device configuration.
AAA Policy Default List	Select this option to apply the default authorization list defined in the device's AAA policy to the EXEC commands associated with this privilege level. See Command Accounting Dialog Box, page 63-13.
Custom Method List	Select this option to define an authorization method list for this privilege level.
Prioritized Method List	Applies only when the Custom Method List option is selected. Defines a sequential list of methods to be queried when authorizing a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to authorize users using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.

Command Accounting Dialog Box—Line Access

Use the Command Accounting dialog box to define which methods to use when recording information about the EXEC commands that are executed for a given privilege. Each accounting record includes a list of the commands executed for that privilege level, as well as the date and time each command was executed, and the name of the user who executed it.

Navigation Path

From the Console Page—Accounting Tab, page 63-47 or the VTY Line Dialog Box—Accounting Tab, page 63-57, click the **Add** button beneath the Command Accounting table.

Related Topics

- Console Policy Page, page 63-42

- VTY Policy Page, page 63-50

Field Reference

Table 63-27 Command Accounting Dialog Box—Line Access

Element	Description
Privilege Level	The privilege level for which you want to define a command accounting list. Valid values range from 0 to 15. **Note** If you do not define a value, level 1 is assigned by default. This value does not appear in the device configuration.
AAA Policy Default List	Select this option to apply the default accounting list defined in the device's AAA policy to the EXEC commands executed for this privilege level.
Custom Method List	Select this option to define an accounting method list for this privilege level.
Generate Accounting Records for	Applies only when Custom Method List is selected. Defines when the device sends an accounting notice to the accounting server: • Start and Stop—Generates accounting records at the beginning and the end of the user process. The user process begins regardless of whether the accounting server receives the "start" accounting record. This is the default. • Stop Only—Generates an accounting record at the end of the user process only. • None—No accounting records are generated.
Prioritized Method List	Applies only when the Custom Method List option is selected. Defines a sequential list of accounting methods to be used when creating accounting records for a user. Enter the names of one or more AAA server group objects (up to four), or click **Select** to select them. Use the up and down arrows in the object selector to define the order in which the selected server groups should be used. If the object that you want is not listed, click the **Create** button to create it. The device tries initially to perform accounting using the first method in the list. If that method fails to respond, the device tries the next method, and so on, until a response is received. **Note** If you select None as a method, it must appear as the last method in the list.
Enable Broadcast to Multiple Servers	Applies only when Custom Method List is selected. When selected, enables the sending of accounting records to multiple AAA servers. Accounting records are sent simultaneously to the first server in each AAA server group defined in the method list. If the first server is unavailable, failover occurs using the backup servers defined within that group. When deselected, accounting records are sent only to the first server in the first AAA server group defined in the method list.

Optional SSH Settings on Cisco IOS Routers

Secure Shell (SSH) is an application and a protocol that uses encryption to provide secure communication between a client and server. You can use SSH to connect remotely to a Cisco IOS router over a VTY line and establish an EXEC session. SSH is the recommended replacement for other protocols, such as Telnet and rlogin, in environments where security is a concern.

All Cisco IOS routers are required to have SSH configured before they can be added to Security Manager. This is because Security Manager uses SSH (in addition to SSL) to communicate with them. The SSH policy provides a way to modify selected default settings and configure selected optional settings.

Related Topics

- Defining Optional SSH Settings, page 63-63
- Chapter 2, "Preparing Devices for Management"
- Setting Up SSH, page 2-5

Defining Optional SSH Settings

SSH is configured by default with the following settings:

- Both SSH version 1 and SSH version 2 are supported.
- The negotiation phase is terminated if not completed successfully after 120 seconds.
- The router tries 3 times to authenticate SSH clients before disconnecting.

You can use Security Manager to modify these default settings and optionally configure the following settings:

- The source interface for SSH packets.
- The name of the RSA key pair to use.
- Whether to regenerate the key during the next deployment.

Before You Begin

- Make sure that SSH is enabled on the router. See Chapter 2, "Preparing Devices for Management".
- Make sure that the VTY lines on the router allow inbound SSH traffic. See Defining VTY Line Setup Parameters, page 63-38.
- Make sure that a hostname and domain name are configured on the router (unless you plan to use a different RSA key pair). You can use the CLI or the Hostname policy in Security Manager for this purpose. See Hostnames and Domain Names on Cisco IOS Routers, page 63-77.

Related Topics

- Optional SSH Settings on Cisco IOS Routers, page 63-63
- Setting Up SSH, page 2-5

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > Secure Shell** from the Policy selector.

- (Policy view) Select **Router Platform > Device Admin > Device Access > Secure Shell** from the Policy Type selector. Select an existing policy or create a new one.

The Secure Shell page is displayed. See Table 63-28 on page 63-65 for a description of the fields on this page.

Step 2 (Optional) Modify the following default settings:

a. The version of SSH to support.

b. The timeout for completing the negotiation phase of the SSH connection.

c. The number of times to attempt authentication of the SSH client.

Step 3 (Optional) In the Source Interface field, enter the name of the interface or interface role whose address should be used as the source interface for all SSH packets sent to SSH clients, or click **Select** to select an interface role object from a list or to create a new one. The source interface must have an IP address.

If you do not enter a value in this field, the address of the closest interface to the destination is used.

Step 4 (Optional) Enter the name of the RSA key pair to use for SSH connections. If you do not enter a value in this field, the router uses the key pair that is based on the hostname and domain name.

Tip Use the CLI command show crypto key mypubkey rsa to display the names and values of each key pair configured on the device.

Step 5 (Optional) Select the **Regenerate Key During Deployment** check box if you want the router to regenerate the RSA key pair used for SSH. This option is useful if you believe that the secrecy of the keys might be compromised. Enter the size of the modulus to use to regenerate the keys.

Note You must remember to return to this policy after deployment to deselect the check box. If you do not do this, a new key is generated during each deployment.

Note This option requires interaction with the device during deployment. Therefore, you should use it only when deploying to live devices, not when deploying to a file.

Note A key pair must already exist on the device *before* you select this option; otherwise, deployment will fail. (This will typically be the case, since IOS routers must have SSH enabled to be added to Security Manager.)

Secure Shell Policy Page

Use the Secure Shell page to change the default SSH settings on the router and to define additional optional settings, if required.

For more information, see Optional SSH Settings on Cisco IOS Routers, page 63-63.

Note You must configure SSH on the device using CLI commands before adding the device to Security Manager. This is because Security Manager uses SSH (as well as SSL) to communicate with Cisco IOS routers. For more information, see Setting Up SSH, page 2-5.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > Secure Shell** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > Secure Shell** from the Policy Type selector. Right-click **Secure Shell** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Chapter 2, "Preparing Devices for Management"
- VTY Policy Page, page 63-50
- Console Policy Page, page 63-42

Field Reference

Table 63-28 Secure Shell Page

Element	Description
SSH Version	The version of SSH to use when connecting to the router: • 1 and 2—SSH version 1 and SSH version 2. This is the default. • 1—SSH version 1 only. • 2—SSH version 2 only.
Timeout	The amount of time the router should wait for the SSH client to respond during the negotiation phase before disconnecting. The default value (and the maximum) is 120 seconds. **Note** After negotiation finishes and the EXEC session begins, the timeout configured for the VTY line applies. See VTY Line Dialog Box—Setup Tab, page 63-52.
Authentication Retries	The number of times the router attempts to authenticate SSH clients. Valid values range from 0 to 5. The default is 3.
Source Interface	The source address for all SSH packets sent to the SSH client. If you do not define a value in this field, the address of the closest interface to the destination (that is, the output interface through which SSH packets are sent) is used. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one.

Table 63-28 Secure Shell Page (continued)

Element	Description
RSA Key Pair	The name of the RSA key pair to use for SSH connections. If you do not enter a value, the router uses the RSA key pair generated from its hostname and domain name. This is the default. **Tip** Use the CLI command **show crypto key mypubkey rsa** to display the names and values of each key pair configured on the device. These are the valid names that can be entered in this field.
Regenerate Key During Deployment	When selected, regenerates the RSA key pair on the router during the next deployment. This option is useful if you are concerned that the secrecy of the keys might be compromised. When deselected, a new key pair is not generated. **Note** This check box is *not* deselected automatically after deployment. If you do not return to this policy to deselect the check box, the key is regenerated each time you deploy. **Note** This option requires interaction with the device during deployment. Therefore, you should use it only when deploying to live devices, not when deploying to a file. **Note** A key pair must already exist on the device *before* you select this option; otherwise, deployment will fail. (This will typically be the case, since IOS routers must have SSH enabled in order to be added to Security Manager.)
Modulus Size	Applies only when the Regenerate Key check box is selected. The size of the modulus used to generate a new key pair. A larger modulus is more secure but takes longer to generate. Valid values range from 360 to 2048 bits. The default is 1024 bits.

SNMP on Cisco IOS Routers

Simple Network Management Protocol (SNMP) defines a standard way for network management stations or workstations to monitor the health and status of many types of devices, including switches, routers, and firewall devices. It comprises a protocol, a database-structure specification, and a set of management data objects. Each SNMP device or member is part of a *community*, which determines the access that each device has (read-only or read-write).

SNMP obtains information from the managed device through a Management Information Base (MIB). The MIB is a database of code blocks called MIB objects, each of which controls one specific function. The MIB object comprises MIB variables, which define the MIB object name, description, default value, and so forth. MIB objects are structured hierarchically in a MIB tree.

SNMP policies enable you to configure the behavior of the SNMP agent running on the router. The agent sends unsolicited information back to the SNMP host as events occur. These unsolicited messages, which are generated in response to significant, predetermined events on the router, are called traps.

The following topics describe the tasks you perform to create SNMP policies on Cisco IOS routers:

- Defining SNMP Agent Properties, page 63-67

- Enabling SNMP Traps, page 63-68

Defining SNMP Agent Properties

When you define the properties of the SNMP agent, you must define the community string and community string type, as well as the address and properties of the SNMP host that receives the traps.

SNMP community strings are embedded passwords to MIBs, which store data about the router's operation and are meant to be available to authenticated remote users. Two types of community strings exist: "public" community strings, which provide read-only access to all objects in the MIB (except community strings themselves), and "private" community strings, which provide read-write access to all objects in the MIB (except community strings).

SNMP hosts receive the traps generated by the router. You must define the address, password, and port number for accessing the SNMP host, as well as the SNMP version being used. Security Manager supports SNMP version 1, version 2c (also called "community-based SNMP") and version 3, which offers authentication and encryption.

Related Topics

- SNMP on Cisco IOS Routers, page 63-66

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > SNMP** from the Policy Type selector. Select an existing policy or create a new one.

The SNMP page is displayed. See Table 63-29 on page 63-69 for a description of the fields on this page.

Step 2 Define the community string needed to access the MIB:

a. Under Permissions, click **Add** to display the Permission dialog box.

b. Define the string. See Table 63-30 on page 63-70 for a description of the available fields.

c. Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the Permissions table.

Note A warning is displayed if you attempt to edit or delete a community string that is in use by an SNMP host. If you continue with the operation, the device creates a private, read-only string that matches the definition for the host in the Trap Receiver table.

Step 3 Define the SNMP host that receives the traps generated by the SNMP agent:

a. Under Trap Receiver, click **Add** to display the Trap Receiver dialog box.

b. Define the host. See Table 63-31 on page 63-71 for a description of the available fields.

c. Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the Trap Receiver table.

Step 4 Under SNMP Server Properties, enter the location and contact information for the administrator responsible for routers configured with this SNMP policy.

This definition, which is text-only and does not affect the operation of the router, provides useful information to the manager of the SNMP host when the manager investigates a particular trap.

Step 5 Click **Configure Traps** to display the SNMP Traps dialog box, which is used to select which traps to enable on the router. For more information, see Enabling SNMP Traps, page 63-68.

Enabling SNMP Traps

The router immediately sends notifications, also called SNMP traps, to the designated SNMP host (management station) when a defined condition occurs, such as a link up, link down, or a syslog event.

To enable SNMP traps, select the check box next to each relevant trap. Certain check boxes activate multiple, related traps.

Note Each trap that you enable consumes system resources. To lessen the impact on system performance, select only those traps that you need for network monitoring.

Related Topics

- SNMP on Cisco IOS Routers, page 63-66

Step 1 Open the SNMP page for defining SNMP server policies on Cisco IOS routers, as described in Defining SNMP Agent Properties, page 63-67.

Step 2 On the SNMP page, click **Configure Traps**. The SNMP Traps dialog box is displayed.

Step 3 Select the check box next to each type of trap to enable. The traps are divided into the following four categories:

- Standard SNMP traps (for example, Authentication, Cold Start, and Warm Start).
- ISAKMP traps (related to Phase 1 of the IPsec process).
- IPsec traps (related to Phase 2 of the IPsec process).
- Other traps (includes syslog messages, protocol-related notifications, and CPU usage warnings).

See Table 63-32 on page 63-73 for a description of the available traps.

Note You must add command-line interface (CLI) commands to fully implement the IP multicast and CPU traps. One method available for entering these commands is by using FlexConfigs. See Chapter 7, “Managing FlexConfigs”.

Tip Click **Select All** to enable all traps displayed in the dialog box or **Deselect All** to disable all the traps.

Step 4 Click **OK** to save your definitions locally on the client and close the dialog box.

Tip To configure SNMP traps not included in this dialog box, define a FlexConfig. See Chapter 7, “Managing FlexConfigs” for more information.

SNMP Policy Page

Use the SNMP page to configure the parameters necessary to send traps from the router to a designated SNMP host. These traps are unsolicited messages that notify the SNMP host of important events occurring on the router.

For more information, see Defining SNMP Agent Properties, page 63-67.

Navigation Path

- (Device view) Select **Platform > Device Admin > Device Access > SNMP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Device Access > SNMP** from the Policy Type selector. Right-click **SNMP** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- SNMP on Cisco IOS Routers, page 63-66
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-29 SNMP Page

Element	Description
Permissions table	
Community String	The community string used for accessing the router's MIB.
Type	The community string type—read-only or read-write.
ACL	The standard ACL that defines the IP addresses permitted to access the router's MIB.
Add button	Opens the Permission Dialog Box, page 63-70. From here you can enter the community string and type required to generate traps.
Edit button	Opens the Permission Dialog Box, page 63-70. From here you can edit the selected permissions profile.
Delete button	Deletes the selected permissions profiles from the table.
Trap Receiver table	
Host IP Address	The IP address of the SNMP host receiving the traps generated by the router.
SNMP Version	The SNMP version being used by the router.
UDP Port	The UDP port that is being used by the SNMP host.
Add button	Opens the Trap Receiver Dialog Box, page 63-71. From here you can define the SNMP host that receives the traps generated by the router.
Edit button	Open the Trap Receiver Dialog Box, page 63-71. From here you can edit the selected SNMP host.
Delete button	Deletes the selected SNMP hosts from the table.
Additional fields and buttons	

Table 63-29 SNMP Page (continued)

Element	Description
SNMP Server Properties	The name and contact information of the system administrator responsible for the SNMP server/agent (that is, the router). The person managing the SNMP host can use this information when tracking down the source of unusual events. The maximum length of each of these properties is 255 characters, including spaces. **Note** The values entered in these fields are text-only and do not affect the operation of the router.
Configure Traps button	Opens a dialog box for selecting which SNMP traps the router should generate. See SNMP Traps Dialog Box, page 63-72.

Permission Dialog Box

Use the Permission dialog box to define the community string and string type required by the SNMP policy. The community string is an embedded password for accessing the Management Information Base (MIB) that stores operational data about the router.

Navigation Path

Go to SNMP Policy Page, page 63-69, then click the **Add** or **Edit** button beneath the Permissions table.

Related Topics

- SNMP Policy Page, page 63-69
- Trap Receiver Dialog Box, page 63-71
- SNMP Traps Dialog Box, page 63-72
- Defining SNMP Agent Properties, page 63-67
- SNMP on Cisco IOS Routers, page 63-66

Field Reference

Table 63-30 Permission Dialog Box

Element	Description
Community String	The community string for accessing the router's MIB. String length ranges from 1 to 128 characters.
Access Control Lists	Applies only to routers running Cisco IOS Software Release 12.3(2)T and up (T-train) or any 12.4 version. The standard ACL containing the IP addresses that can access the router's MIB. Defining an ACL provides an additional layer of security by limiting the source addresses that can make use of the community string. Enter the name of a standard ACL object, or click **Select** to select the object from a list or to create a new one.

Table 63-30 Permission Dialog Box (continued)

Element	Description
Read-Write	This community string type provides read-write access to all objects in the MIB (except community strings).
Read-Only	This community string type provides read-only access to all objects in the MIB (except community strings). This is the default.

Trap Receiver Dialog Box

Use the Trap Receiver dialog box to define the SNMP hosts that receive traps generated by the router. This includes defining the version of SNMP to use.

Navigation Path

Go to the SNMP Policy Page, page 63-69, then click the **Add** or **Edit** button beneath the Trap Receiver table.

Related Topics

- SNMP Policy Page, page 63-69
- Permission Dialog Box, page 63-70
- SNMP Traps Dialog Box, page 63-72
- Defining SNMP Agent Properties, page 63-67
- SNMP on Cisco IOS Routers, page 63-66

Field Reference

Table 63-31 Trap Receiver Dialog Box

Element	Description
Host IP Address	The IP address of the SNMP host receiving the traps generated by the router. Enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one.
SNMP Version	The version of SNMP to use—version 1, version 2c, or version 3.
Community String	Applies only when version 1 or version 2c is selected. The password required to access the SNMP host. Enter the string again in the Confirm field. **Note** We recommend that you use one of the strings defined in the Permissions table as the password to the SNMP host. You may, however, enter a different password. String length ranges from 1 to 128 characters. Your entry does not appear in the Permissions table and is read-only.

Table 63-31 Trap Receiver Dialog Box (continued)

Element	Description
User Name	Applies only when version 3 is selected. The password required to access the SNMP host. Enter the string again in the Confirm field. **Note** We recommend that you use one of the strings defined in the Permissions table as the password to the SNMP host. You may, however, enter a different password. String length ranges from 1 to 128 characters. Your entry does not appear in the Permissions table and is read-only.
SNMPv3 Security	Applies only when version 3 is selected. The level of security to apply to SNMP traffic: • No MD5, No DES—No packet authentication. • MD5 (auth)—MD5 authentication, but no encryption. • DES (priv)—MD5 authentication and DES encryption.
UDP Port	The port number for the SNMP host. The default is 162. Valid values range from 0 to 65535.

SNMP Traps Dialog Box

Use the SNMP Traps dialog box to select the events in the router that should generate SNMP traps. To lessen possible degradation of system performance, select only those traps that are needed for network monitoring purposes.

Tip You can configure SNMP traps not included in this dialog box by defining FlexConfigs. For more information, see Understanding FlexConfig Policies and Policy Objects, page 7-2.

Navigation Path

Go to the SNMP Policy Page, page 63-69, then click **Configure Traps**.

Related Topics

- SNMP Policy Page, page 63-69
- Permission Dialog Box, page 63-70
- Trap Receiver Dialog Box, page 63-71
- Enabling SNMP Traps, page 63-68
- SNMP on Cisco IOS Routers, page 63-66

Field Reference

Table 63-32 SNMP Traps Dialog Box

Element	Description
Standard SNMP Traps	Enables or disables standard SNMP traps. Options are: • Cold start—Sends a trap when the router reinitializes in a way that could change the configuration of the SNMP agent (or any other trap-receiving entity). • Warm start—Sends a trap when the router reinitializes in a way that does not change the configuration of the SNMP agent (or any other trap-receiving entity). • Authentication—Sends a trap if an SNMP request from the SNMP host fails because of an invalid community string.
IPsec Traps	Enables or disables individual IPsec-related traps. Options are: • Cryptomap—Sends a trap when a crypto map entry is added to, or removed from, the device's crypto map set. Additionally, this option sends a trap when a crypto map set is attached to, or detached from, an active interface. • Too Many SAs—Sends a trap if an attempt is made to create a security association (SA) when there is insufficient memory on the device. • Tunnel—Sends a trap when an IPsec Phase 2 tunnel becomes active or inactive. For more information, see Understanding IPsec Proposals for Site-to-Site VPNs, page 26-19.
ISAKMP Traps	Enables or disables individual Internet Security Association and Key Exchange Protocol (ISAKMP) traps. Options are: • Policy—Sends a trap when an ISAKMP policy is created or deleted. • Tunnel—Sends a trap when a Phase 1 IKE tunnel becomes active or inactive. For more information, see Understanding IKE, page 26-5.

Table 63-32 SNMP Traps Dialog Box (continued)

Element	Description
Other Traps	Enables or disables additional SNMP traps. Options are: • Syslog—Sends syslog messages to the SNMP host. • TTY—Sends Cisco-specific notifications when a Transmission Control Protocol (TCP) connection closes. • BGP—Sends notifications when Border Gateway Protocol (BGP) state changes occur. See BGP Routing on Cisco IOS Routers, page 67-1. • IP Multicast—(Applicable to multicast routers only) Sends a trap if the router fails to receive a defined number of heartbeat packets from heartbeat sources within a defined time interval. If you select IP Multicast, you must also manually configure the **ip multicast heartbeat** command on the device to configure the multicast address and heartbeat limits. You can use FlexConfigs to do this. • CPU—Sends a trap when CPU usage rises and remains above an upper threshold or falls and remains below a lower threshold. If you select CPU, you must also manually configure the **process cpu threshold type** command on the device to configure the thresholds. You can use FlexConfigs to do this. • HSRP—Sends Hot Standby Routing Protocol (HSRP) notifications.
Select All button	Enables all the SNMP traps displayed in the dialog box.
Deselect All button	Disables all the SNMP traps displayed in the dialog box.

DNS on Cisco IOS Routers

The Domain Name System (DNS) is a distributed database in which you can map hostnames to IP addresses through the DNS protocol from a DNS server. Each unique IP address can have an associated hostname. DNS is what makes it possible to connect to hosts without having to know the 32-bit IP address of that host. The DNS server takes the provided hostname and translates it into the appropriate IP address.

In addition to the translation provided by remote DNS servers, you can configure Cisco IOS routers with a local host table containing static mappings of hosts to IP addresses. When commands such as connect, telnet, and ping are used, the router checks this host table before querying the DNS servers, which speeds the translation process.

By default, the DNS feature is enabled on all Cisco IOS routers.

Related Topics

- Defining DNS Policies, page 63-75

Defining DNS Policies

When you define a DNS policy in Security Manager, you can specify the remote DNS servers used by the router for hostname-to-address translations. In addition, you can define a static host table that contains local translations used exclusively by this device. Having selected addresses in this type of cache can speed the translation process by eliminating the need to query the DNS servers.

Related Topics

- DNS on Cisco IOS Routers, page 63-74

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > DNS** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > DNS** from the Policy Type selector. Select an existing policy or create a new one.

The DNS page is displayed. See Table 63-33 on page 63-76 for a description of the fields on this page.

Step 2 In the Servers field, enter the addresses of the DNS servers (up to 6) that can perform hostname-to-address translations for the router. You can use a combination of addresses and network/host objects, or click **Select** to display a selector. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Tip If the network you want is not listed in the selector, click the **Create** button or the **Edit** button in the selector to display the Add or Edit Network/Host Dialog Box, page 6-83. From here you can create a network/host object to use in the policy.

Step 3 (Optional) In the Hosts field, enter the static host mappings that you want to define in the router's host table:

a. Click **Add** to display the IP Host Dialog Box, page 63-76.

b. Enter the hostname to translate.

c. Enter up to three addresses or network/host objects, or click **Select** to display a selector. These are the addresses to which the router translates the hostname.

d. Click **OK**. The mapping is displayed in the Hosts field on the DNS page.

e. Repeat a. through d. to add more hosts to the host table.

Note To edit a host mapping, select the definition from the Hosts field, then click **Edit**. To remove a host mapping, select it, then click **Delete**.

Step 4 (Optional) Deselect the **Domain Lookup** check box to disable DNS functionality on the router.

DNS Policy Page

Use the DNS policy page to define the local IP host table and the Domain Name System (DNS) servers that the router should use for translating hostnames to IP addresses. You can also prevent the router from performing DNS lookups by disabling the DNS feature.

Navigation Path

- (Device view) Select **Platform > Device Admin > DNS** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > DNS** from the Policy Type selector. Right-click **DNS** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- DNS on Cisco IOS Routers, page 63-74

Field Reference

Table 63-33 DNS Page

Element	Description
Servers	The DNS servers used by the router to perform DNS lookups. Enter one or more addresses or network/host objects, or click **Select** to select an object from a list or to create a new one. You can define a maximum of six DNS servers.
Hosts	The local host table configured on the router. When a user types in a hostname, the router checks this table first before querying the DNS servers defined in the Servers field. Click **Add** to display the IP Host Dialog Box, page 63-76. From here you can define a hostname and the IP addresses to associate with that hostname. **Note** To edit an entry in the host table, select it, then click **Edit**. To remove an entry, select it, then click **Delete**.
Domain Lookup	When selected, the router performs lookups on the defined DNS servers. This is the default. When deselected, lookups on remote DNS servers are disabled.

IP Host Dialog Box

Use the IP Host dialog box to configure the host table on the router. This is the table of static, local mappings that the router uses to translate hostnames to IP addresses. If the router does not find the required entry in the host table, it queries the DNS servers that are defined on the DNS page.

Navigation Path

Go to the DNS Policy Page, page 63-76, then click **Add** under Hosts.

Related Topics

- DNS on Cisco IOS Routers, page 63-74

Field Reference

Table 63-34 IP Host Dialog Box

Element	Description
Host Name	The hostname to include in the router's local host table.
Addresses	The addresses to associate with the hostname. Enter one or more addresses or network/host objects, or click **Select** to select an object from a list or to create a new object. You can define a maximum of three addresses per hostname.

Hostnames and Domain Names on Cisco IOS Routers

The Hostname policy configures the hostname and domain name of the selected router. After you deploy this policy, any changes that you made to the hostname and domain name are reflected in the Device Properties page (see Viewing or Changing Device Properties, page 3-41).

Related Topics

- Defining Hostname Policies, page 63-77

Defining Hostname Policies

When you define a hostname policy, Security Manager updates the hostname and domain name fields in the Device Properties dialog box after deployment. See Viewing or Changing Device Properties, page 3-41.

Related Topics

- Hostnames and Domain Names on Cisco IOS Routers, page 63-77

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Hostname** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Hostname** from the Policy Type selector. Select an existing policy or create a new one.

The Hostname page is displayed. See Table 63-35 on page 63-78 for a description of the fields on this page.

Step 2 Enter the hostname for the router. Names must start with a letter, end with a letter or digit, and include only letters, digits, and hyphens. The maximum length is 63 characters.

Step 3 Enter the domain name for the router. The router uses this domain name for RSA key generation and in policies when you do not enter the fully-qualified domain name.

Hostname Policy Page

Use the Hostname page to define the hostname and domain name assigned to the router. For more information, see Defining Hostname Policies, page 63-77.

Navigation Path

- (Device view) Select **Platform > Device Admin > Hostname** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Hostname** from the Policy Type selector. Right-click **Hostname** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Hostnames and Domain Names on Cisco IOS Routers, page 63-77

Field Reference

Table 63-35 Hostname Page

Element	Description
Host Name	The hostname of the router. Names must start with a letter, end with a letter or digit, and include only letters, digits, and hyphens. The maximum length is 63 characters.
Domain Name	The default domain name of the router. The maximum length is 63 characters. The router uses this domain name for RSA key generation and in policies when you do not enter the fully-qualified domain name (FQDN).

Memory Settings on Cisco IOS Routers

The Memory policy configures settings relating to router memory. This policy provides you with methods for monitoring memory consumption, including the ability to generate notification messages when available memory drops below predefined thresholds.

Note The Memory policy is supported on routers running Cisco IOS Software Release 12.3(14)T or later.

Related Topics

- Defining Router Memory Settings, page 63-78

Defining Router Memory Settings

You can use Security Manager to modify the following default memory settings:

- The number of hours that the router maintains the log of memory consumption.
- Whether to enable the Memory Allocation Lite feature.
- The amount of memory to reserve for critical system log messages.

In addition, you can define:

- The lower thresholds for processor and I/O memory. Log messages are sent when available memory drops below these thresholds.
- The types of sanity checks to perform.

Related Topics

- Memory Settings on Cisco IOS Routers, page 63-78
- Logging on Cisco IOS Routers, page 65-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Memory** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Memory** from the Policy Type selector. Select an existing policy or create a new one.

The Memory page is displayed.

Step 2 (Optional) Define the memory settings of the router, as required. See Table 63-36 on page 63-80 for a description of the available fields.

Memory Policy Page

Use the Memory page to define settings related to router memory, including:

- The amount of time to retain the memory log.
- The thresholds for available processor and I/O memory.
- The amount of memory reserved for critical log messages.
- Whether to perform sanity checks on buffers and queues.
- Whether to enable the “memory-allocation lite” feature.

For more information, see Defining Router Memory Settings, page 63-78.

Navigation Path

- (Device view) Select **Platform > Device Admin > Memory** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Memory** from the Policy Type selector. Right-click **Memory** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Memory Settings on Cisco IOS Routers, page 63-78
- CPU Policy Page, page 63-26
- Syslog Logging Setup Policy Page, page 65-7
- Syslog Servers Policy Page, page 65-10

Field Reference

Table 63-36 Memory Page

Element	Description
Maintain Memory Log	The number of hours that the router should maintain the log containing the history of memory consumption on the device. Valid values range from 12 to 72 hours. The default is 24 (1 day). **Note** The memory log is enabled by default and cannot be disabled.
Processor Threshold	The processor memory threshold in kilobytes. When available processor memory falls below this threshold, a notification message is triggered. Valid values range from 1 to 4294967295 kilobytes (4096 gigabytes). **Note** Another notification message is generated when available free memory rises to 5% above the threshold.
I/O Threshold	The I/O memory threshold in kilobytes. When available processor memory falls below this threshold, a notification message is triggered. Valid values range from 1 to 4294967295 kilobytes (4096 gigabytes). **Note** Another notification message is generated when available free memory rises to 5% above the threshold.
Memory Allocation Lite	When selected, the "memory-allocation lite" (malloc_lite) feature on the router is enabled. This feature avoids excessive memory allocation overhead for situations where less than 128 bytes are required. This is the default. When deselected, the "memory-allocation lite" feature is disabled. **Note** This feature is supported for processor memory pools only.
Memory Region For Critical Notifications	The amount of memory (in kilobytes) reserved for critical system log messages. Valid values range from 1 to 4294967295 kilobytes (4096 gigabytes), but the value you specify cannot exceed 25% of total memory. This option reserves a region of memory on the router so that the router can issue critical system log messages even when system resources are overloaded.
Perform Sanity Checks	The types of sanity checks to perform: • Buffer—When selected, performs sanity checks on all buffers. Sanity checks are performed when a packet buffer is allocated and when the packet buffer is returned to the buffer pool. • Queue—When selected, performs sanity checks on all queues. • All—When selected, performs sanity checks on all buffers and queues. **Note** Enabling any of these options may result in a slight degradation of router performance.

Secure Device Provisioning on Cisco IOS Routers

Secure Device Provisioning (SDP) offers an integrated solution for streamlining VPN and network security deployment. SDP (previously called Easy Secure Device Deployment, or EzSDD) enables remote-site users to securely bootstrap their VPN device through an easy-to-use web interface, thereby easing the deployment burden, lowering costs, and shortening the network development cycle. For example, a telecommuter or small branch office user can remove a new device from its shipping package, plug it in, open a simple web management interface, and establish VPN connectivity, all within a period of just a few minutes.

For more information about SDP, see *Setting Up Secure Device Provisioning (SDP) for Enrollment in a PKI*, which can be found in *Cisco IOS Security Configuration Guide, Release 12.4T*.

Note SDP requires Cisco IOS Software Release 12.3(8)T or later. Attempting to deploy this policy to a router running an earlier version could result in deployment failure. You also cannot configure the policy on NO-VPN router models (those that do not allow VPN configurations, such as the 3845 NOVPN).

Trusted Transitive Introduction (TTI) is the protocol that acts as the primary mechanism for implementing SDP. As shown in Figure 63-3, TTI comprises the following three entities:

- Introducer—A mutually trusted device that introduces the petitioner to the registrar. Introducers can be end users who use SDP to deploy VPN devices associated with themselves to the PKI network, or an administrator/management system that uses SDP to deploy many VPN devices to the PKI network. This latter type is known as an administrative introducer. For more information, see Configuring a AAA Server Group for Administrative Introducers, page 63-84.
- Petitioner—A remote-site device that is joined to the secure domain. The petitioner serves web pages to the introducer and receives the bootstrap configuration from the introducer's web browser. The petitioner component is enabled by default on all Cisco IOS devices.
- Registrar—A server that authorizes the petitioner by communicating directly with an authentication, authorization, and accounting (AAA) server to verify user credentials, permit or deny enrollment, and retrieve user-specific configuration information.

Use the SDP policy in Security Manager to configure the router as a registrar.

Figure 63-3 Secure Device Provisioning

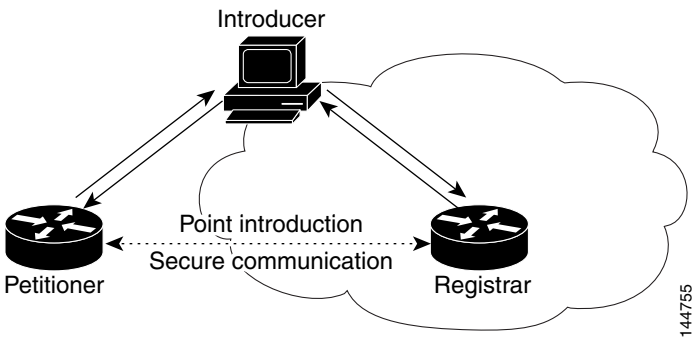

For more information about Secure Device Provisioning, see:

- Contents of Bootstrap Configuration, page 63-82
- Secure Device Provisioning Workflow, page 63-82

- Defining Secure Device Provisioning Policies, page 63-83

Contents of Bootstrap Configuration

The bootstrap configuration provided by SDP typically does the following:

- Sets the petitioner's hostname.
- Synchronizes the petitioner's system clock with the registrar.
- Sets the petitioner's trustpoint.
- Sets the petitioner's authentication and authorization mechanism.
- Pushes the CA certificate.
- Enrolls the petitioner with the PKI server.
- Sets other VPN configurations, such as the configuration required to establish a management tunnel.
- Sets Cisco Networking Services (CNS) configuration.
- Sets the petitioner's DHCP pool.

Related Topics

- Secure Device Provisioning Workflow, page 63-82
- Secure Device Provisioning on Cisco IOS Routers, page 63-81

Secure Device Provisioning Workflow

The following illustrates the steps required to use SDP to register a remote-site device in a secure network:

1. Unpack the router and connect the power, LAN, and WAN cables.
2. Turn on a computer (introducer) that is assigned an IP address from the DHCP server on the router, open a web browser, and go to the petitioner URL (http://device/ezsdd/welcome) on the router. The router responds with a registration page (also called the local login dialog box).
3. Enter the username and password, then click **OK**. On the welcome page, enter the URL for the registrar. The following actions occur:
   a. The browser opens an HTTPS-secured session to the central-site registrar, which verifies the username with the AAA server and returns the appropriate bootstrap configuration to the browser.
   b. The browser feeds the bootstrap configuration to the remote-site router, configuring PKI trustpoint enrollment and IPsec VPN connectivity, and provisioning system attributes and other information.
   c. You are notified that bootstrap configuration is complete.

   Related Topics

   - Contents of Bootstrap Configuration, page 63-82
   - Secure Device Provisioning on Cisco IOS Routers, page 63-81

Defining Secure Device Provisioning Policies

The petitioner component is automatically enabled on all Cisco IOS routers. The SDP policy in Security Manager enables the registrar. To define an SDP policy you must define:

- The AAA server group containing the AAA server that the registrar uses to authenticate and authorize the introducer.
- The CA server to which the petitioner enrolls during the bootstrap process.
- The location of the introduction page that is displayed after authorization was performed.
- The location of the bootstrap configuration to be provided to the petitioner.

Related Topics

- Secure Device Provisioning Workflow, page 63-82
- Configuring a AAA Server Group for Administrative Introducers, page 63-84
- Secure Device Provisioning on Cisco IOS Routers, page 63-81

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Secure Device Provisioning** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Secure Device Provisioning** from the Policy Type selector. Select an existing policy or create a new one.

The Secure Device Provisioning page is displayed. See Table 63-37 on page 63-85 for a description of the fields on this page.

Step 2 Under Introducer Authentication, enter the name of the AAA server group containing the relevant AAA server, or click **Select** to select it from a list or to create a new object.

The selected AAA server determines whether the username and password supplied by the introducer represent an authorized user. The AAA server must use TACACS+, RADIUS, or be local.

Note Each AAA server in the selected group must be configured to communicate with an interface that exists on the router; otherwise, validation fails. If you want to configure a different AAA server group for authenticating and authorizing administrative introducers, see Configuring a AAA Server Group for Administrative Introducers, page 63-84.

Step 3 Under Petitioner Authentication, define the CA server that authenticates the identity of the petitioner by doing one of the following:

- Select **Local CA Server**, then enter the local CA name in the field provided. If you have already configured the CA server locally on the registrar, a trustpoint is generated automatically.

  Note If you have not configured the router as the CA server, enter the command **Crypto pki server [name]** using the CLI or FlexConfigs. This command is mandatory when you deploy an SDP policy configured with a local CA server.

- Select Remote CA Server, then enter the name of a PKI enrollment object, or click **Select** to select it from a list or to create a new object.

  The PKI enrollment object defines the external CA server used in the SDP policy.

Step 4 Select the source of the introduction page that is displayed after you log in to the registrar. The introduction page indicates whether authorization was successfully completed and contains a button for completing the process of obtaining the bootstrap configuration.

If you do not select the default welcome page, you must enter the URL required to access a different welcome page that you prepared elsewhere.

Step 5 Select the source of the bootstrap configuration provided to the petitioner to implement its first-time configuration:

- If the source of the bootstrap configuration is a non-Security Manager URL, enter the URL and also the username and password for accessing the URL, if required.
- If the source of the configuration file is a Security Manager URL:
  - Enter the name of a FlexConfig, or click **Select** to select it from a list or to create a new object. The FlexConfig contains the device commands required to retrieve the appropriate bootstrap configuration. For more information, see Add or Edit FlexConfig Dialog Box, page 7-30.
  - Enter the device name formula required by the FlexConfig to derive the device name of the petitioner from the username submitted by the introducer. (The two names typically have a fixed relationship.) The default formula is $n, which uses the introducer name to determine the device name.

    The device name determines which bootstrap configuration the petitioner should receive. The resulting URL contains the name of the FlexConfig you selected, as well as the parameters and formula you defined.
  - Enter a username and password for accessing the Security Manager server containing the FlexConfig. The password can contain alphanumeric characters, but cannot consist of a single digit.

Configuring a AAA Server Group for Administrative Introducers

Administrative introducers are administrators or management systems that introduce many devices to the PKI network. You can configure a AAA server group for authenticating and authorizing administrative introducers by appending the following FlexConfig to the configuration of the router:

```
aaa new-model
radius-server host 1.2.3.4 auth-port 1645 acct-port 1646 key key
aaa group server radius default-radius-group2
server 1.2.3.4 auth-port 1645 acct-port 1646
exit
aaa authentication login CSM_SDP2 group default-radius-group2
crypto provisioning registrar
administrator authentication list CSM_SDP2
administrator authorization list CSM_SDP2
exit
```

This FlexConfig serves two functions—it configures the AAA server group to use and it associates this server group with the SDP crypto.

For more information about administrative introducers, see *Administrative Secure Device Provisioning Introducer* on Cisco.com at this URL:
http://www.cisco.com/en/US/docs/ios/12_3t/12_3t14/feature/guide/gtadintr.html

Related Topics

- Secure Device Provisioning on Cisco IOS Routers, page 63-81
- Defining Secure Device Provisioning Policies, page 63-83
- Understanding FlexConfig Policies and Policy Objects, page 7-2

Secure Device Provisioning Policy Page

Secure Device Provisioning (SDP) policies (formerly known as Easy Secure Device Deployment or EzSDD) enable you to configure a Cisco IOS router as a *registrar*. This is the SDP component that retrieves bootstrap configurations for *petitioners*, which are remote-site devices that are enrolling in the network security infrastructure. These devices uses the bootstrap configuration for first-time configuration purposes. The registrar also verifies the identity of the *introducer*, which is the user who introduces the petitioner to the registrar.

For more information, see Defining Secure Device Provisioning Policies, page 63-83.

Navigation Path

- (Device view) Select **Platform > Device Admin > Secure Device Provisioning** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Secure Device Provisioning** from the Policy Type selector. Create a new policy or select an existing one.

Related Topics

- Secure Device Provisioning on Cisco IOS Routers, page 63-81
- Secure Device Provisioning Workflow, page 63-82
- Understanding AAA Server and Server Group Objects, page 6-27
- Understanding FlexConfig Policies and Policy Objects, page 7-2

Field Reference

Table 63-37 Secure Device Provisioning Page

Element	Description
Introducer Authentication (AAA)	The AAA server group that authenticates the username and password supplied by the introducer. Enter the name of a AAA server group object, or click **Select** to select it from a list or to create a new object. **Note** To configure a separate AAA server group for authenticating administrative introducers, see Configuring a AAA Server Group for Administrative Introducers, page 63-84.

Table 63-37 Secure Device Provisioning Page (continued)

Element	Description
Petitioner Authentication	The CA server that authenticates the identity of the petitioner: • Local CA Server—Select this option when the router itself is already configured to act as the CA server. Enter the name of the local CA in the field provided. **Note** If you have not configured the router as the CA server, enter the command **Crypto pki server [name]** using the CLI or FlexConfigs. This command is mandatory when you deploy an SDP policy configured with a local CA server. • Remote CA Server—Select this option when using an external CA server. Enter the name of a a PKI enrollment object, or click **Select** to select it from a list or to create a new object. For more information about PKI enrollment objects, see PKI Enrollment Dialog Box, page 26-58.
Introduction Page	The source of the introduction page to display to the introducer after authorization is performed: • Use default introduction page—Uses a default page provided with Security Manager. • Specify introduction page URL—Uses the introduction page specified in the URL field. Supported protocols include: FTP, HTTP, HTTPS, null, NVRAM, RCP, SCP, system, TFTP, Webflash, and XMODEM.

Table 63-37 Secure Device Provisioning Page (continued)

Element	Description
Bootstrap Configuration	The source of the bootstrap configuration to provide to the petitioner for first-time configuration: • Non-Security Manager URL—Used when the bootstrap configuration is located externally to Security Manager. Enter its location in the URL field. If required, enter a username and password to access the server containing the bootstrap configuration. • Security Manager URL—Used when Security Manager is providing the bootstrap configuration. Enter information in the following fields: – FlexConfig—The FlexConfig that contains the basic CLI structure required to create the bootstrap configuration. Enter the name of a FlexConfig object, or click **Select** to display a selector. After selecting the FlexConfig, you must enter a username and password to access the Security Manager server that contains the FlexConfig. – Device name formula—The formula required by Security Manager to determine the device name of the petitioner from the username that the introducer supplied. Typically a fixed relationship exists between the username and the device name, which enables a formula like this to be established. The default formula is $n, which uses the introducer name to determine the device name. The device name is required to determine the configuration file that the petitioner should receive. If required, enter a username and password to access the server containing the bootstrap configuration. The password can contain alphanumeric characters, but cannot consist of a single digit.

DHCP on Cisco IOS Routers

In Security Manager, certain security features, such as Easy VPN and 802.1x, require Dynamic Host Configuration Protocol (DHCP) client/server configuration. DHCP is widely used in LAN environments to dynamically assign host IP addresses from a centralized server, which significantly reduces the overhead of administering IP addresses.

DHCP servers assign and manage IP addresses from specified address pools within a router to DHCP clients. If the DHCP server cannot satisfy a DHCP request from its own database, it can forward the request to one or more secondary DHCP servers defined by the network administrator.

Security Manager enables you to configure a Cisco IOS device as a DHCP server for clients (hosts) that are connected to the device's inside interface. When you configure a DHCP server, you use IP pools (a range of IP addresses reserved for a DHCP server). The IP pools you select determine the range of IP

addresses the server can use. These addresses are provided to client devices for a defined period of time called a lease. When this lease expires, the address is returned to the address pool, enabling the DHCP server to assign it to a different device.

For more information about DHCP, see:

- Understanding DHCP Database Agents, page 63-88
- Understanding DHCP Relay Agents, page 63-88
- Understanding DHCP Option 82, page 63-89
- Understanding Secured ARP, page 63-89

To configure a DHCP policy, see:

- Defining DHCP Policies, page 63-90
- Defining DHCP Address Pools, page 63-91

Understanding DHCP Database Agents

A DHCP database agent is any external host—for example, an FTP, TFTP, or RCP server—that stores the DHCP bindings database. You can include one or more DHCP database agents in each DHCP policy, as well as configure the interval between database updates to the agent.

Note If you configure an external DHCP database agent, it is not necessary to define IP address pools, but you may do so. For more information about IP address pools, see Defining DHCP Address Pools, page 63-91.

Related Topics

- Understanding DHCP Relay Agents, page 63-88
- Understanding DHCP Option 82, page 63-89
- Understanding Secured ARP, page 63-89
- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87

Understanding DHCP Relay Agents

A DHCP relay agent is any host that forwards DHCP packets between clients and servers when they do not reside on the same physical subnet. Relay agents receive DHCP messages and then generate a new DHCP message to send on another interface. You can configure a reforwarding policy that determines what the DHCP relay agent should do if a forwarded message already contains relay information.

DHCP relay options in Security Manager include:

- Drop—The relay agent discards messages with existing relay information if Option 82 information is also present.
- Keep—The relay agent retains existing relay information.
- Replace—The relay agent overwrites existing information with its own relay information.

For example, you can have the DHCP relay agent replace the forwarded message with a new relay message. Additionally, you can choose whether to have the relay agent check the validity of relay information contained within forwarded BOOTREPLY messages.

Related Topics

- Understanding DHCP Database Agents, page 63-88
- Understanding DHCP Option 82, page 63-89
- Understanding Secured ARP, page 63-89
- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87

Understanding DHCP Option 82

DHCP option 82 enables the DHCP relay agent to include information about itself and its attached client when it forwards requests from a DHCP client to a DHCP server. The DHCP server can use this information to assign IP addresses, perform access control, and set quality of service (QoS) and security policies for each of its subscribers. When the DHCP option 82 feature is enabled, a subscriber is identified by the switch port through which it connects to the networks, instead of by its MAC address. Multiple hosts on the subscriber LAN can be connected to the same port on the access switch and are uniquely identified. Option 82 also enhances security on access switches by providing the ability to use a user's IP address to locate the port on which a user is attached.

Related Topics

- Understanding DHCP Database Agents, page 63-88
- Understanding DHCP Relay Agents, page 63-88
- Understanding Secured ARP, page 63-89
- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87

Understanding Secured ARP

The DHCP Secure IP Address Assignment feature (also called DHCP Authorized ARP) enables you to secure Address Resolution Protocol (ARP) table entries to DHCP leases in the DHCP database. This feature secures and synchronizes the client's MAC address to the DHCP binding, preventing unauthorized clients or hackers from spoofing the DHCP server and taking over a DHCP lease of an authorized client.

When you enable this feature and the DHCP server assigns an IP address to the DHCP client, the DHCP server adds a secure ARP entry to the ARP table with the assigned IP address and the MAC address of the client. These ARP entries cannot be updated by any other dynamic ARP packets, and they exist in the ARP table for as long as the lease is active.

Secure ARP entries can be deleted only by an explicit termination message from the DHCP client or by the DHCP server when the binding expires. To detect when a client has logged out, Secured ARP sends periodic ARP messages to which only authorized users can respond. Unauthorized responses are blocked at the DHCP server, providing an additional level of security.

Note Secured ARP disables dynamic ARP learning on an interface.

Related Topics

- Understanding DHCP Database Agents, page 63-88
- Understanding DHCP Relay Agents, page 63-88
- Understanding DHCP Option 82, page 63-89
- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87

Defining DHCP Policies

When you configure a DHCP policy, you must define the IP address pools for the server to use to provide addresses to DHCP clients. In addition, you can optionally define the following:

- External DHCP database agent.
- IP ranges to exclude from DHCP.
- DHCP relay parameters.

Note When configuring DHCP on a Cisco IOS router, make sure that the router does not contain an access rule denying Bootstrap Protocol (BootP) traffic. Having such a rule blocks DHCP traffic from being transmitted.

Related Topics

- DHCP on Cisco IOS Routers, page 63-87

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Server Access > DHCP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Server Access > DHCP** from the Policy Type selector. Select an existing policy or create a new one.

The DHCP Policy page is displayed. See Table 63-38 on page 63-92 for a description of the fields on this page.

Step 2 (Optional) Under Databases, click the **Add** button to display the DHCP Database Dialog Box, page 63-94. From here you can define external DHCP database agents. For more information, see Understanding DHCP Database Agents, page 63-88.

Step 3 (Optional) Under Excluded IPs, enter the IP addresses or address ranges within a DHCP address pool that should not be made available to DHCP clients. You can use a combination of addresses and network/host objects, or click **Select** to display a selector. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Tip If the network you want is not listed in the selector, click the **Create** button to display the Add or Edit Network/Host Dialog Box, page 6-83. From here you can create a network/host object.

Step 4 Under IP Pools, click the **Add** button to display the IP Pool Dialog Box, page 63-94. From here you can define the address pools to be used by the DHCP server. For more information, see Defining DHCP Address Pools, page 63-91.

Step 5 (Optional) When you use a relay agent to manage requests from DHCP clients located on a different subnet from the DHCP server, define the following DHCP relay options:

a. Select the relay agent information reforwarding policy (Drop, Keep, or Replace). DHCP relay agents implement this policy when they receive messages already containing relay information.

b. Select the **Option** check box to enable the insertion of Option 82 data in requests that the relay agent forwards to the DHCP server.

c. Select the **Check** check box to validate DHCP Option 82 reply packets sent by the DHCP server.

When you enable this option, invalid messages are dropped. Valid messages are stripped of the option-82 field before they are forwarded to the DHCP client. When you disable this option, the option-82 field is removed from the packet without being checked first for validity.

See Understanding DHCP Relay Agents, page 63-88 for more information.

Defining DHCP Address Pools

When you configure a DHCP policy that does not include an external database agent, you must define at least one IP address pool. This pool contains the addresses that the DHCP server can dynamically assign to DHCP clients. Additionally, you can define the following IP pool-specific options:

- The default routers, DNS servers, WINS servers, and domain used by DHCP clients.
- Whether to use the Secured ARP feature.
- Whether to import information regarding IP pool options from a centralized DHCP server.
- The length of the lease.
- The location of the TFTP server that IP telephony devices require to use addresses from this pool.

Related Topics

- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87

Step 1 On the DHCP page, click the **Create** button under IP Pools. The IP Pool dialog box is displayed.

Step 2 Define the address pool. See Table 63-40 on page 63-95 for a description of the available fields.

Step 3 Click **OK** to save your definitions locally on the client and close the dialog box. The IP pool appears in the table displayed under IP Pools on the DHCP page.

Step 4 Repeat Step 1 through Step 3 to define additional address pools, if required.

Note To edit an IP pool, select it from the table, then click the **Edit** button. To delete an IP pool, select it from the table, then click the **Delete** button. You cannot delete a pool whose addresses have been assigned to DHCP clients.

DHCP Policy Page

Use the DHCP policy page to define a DHCP server policy on the selected router. This includes specifying the address pools used by the DHCP server when assigning addresses to requesting clients.

For more information, see Defining DHCP Policies, page 63-90.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > DHCP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Server Access > DHCP** from the Policy Type selector. Right-click **DHCP** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- DHCP on Cisco IOS Routers, page 63-87
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-38 DHCP Policy Page

Element	Description
Databases Table	
Database URL	The URL of the external DHCP database agent.
Timeout	The amount of time to wait (in seconds) for a response from the external DHCP database agent before aborting a database transfer.
Write Delay	The interval (in seconds) between DHCP assignment updates sent to the external DHCP database agent.
Add button	Opens the DHCP Database Dialog Box, page 63-94. From here you can define a DHCP database agent.
Edit button	Opens the DHCP Database Dialog Box, page 63-94. From here you can edit the selected DHCP database agent.
Delete button	Deletes the selected DHCP database agents.
Excluded IPs	
Excluded IPs or IP Ranges	The IP addresses or address ranges to exclude from DHCP. These addresses are not assigned by the DHCP server to DHCP clients requesting addresses. Enter one or more network addresses or network/host objects, or click **Select** to select an object from a list or to create a new object. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.
IP Pools Table	
Name	The name of the IP pool.
Network	The IP address and subnet mask of the IP pool.

Table 63-38 DHCP Policy Page (continued)

Element	Description
Default Router	The IP addresses of the default routers used by DHCP clients.
DNS Server	The IP addresses of the DNS servers used by DHCP clients.
NetBIOS (WINS) Server	The IP addresses of the Windows Internet Naming Service (WINS) servers used by Microsoft DHCP clients.
Domain Name	The domain name for DHCP clients.
Import All	Indicates whether the remote DHCP server imports certain DHCP options from a centralized DHCP server.
Secured ARP	Indicates whether secured ARP is enabled on this IP pool to help prevent IP spoofing by unauthorized users.
Lease	The duration of the lease for each IP address assigned by the DHCP server from this IP pool.
Option 150	The IP address of the TFTP server required by IP phones for configuration, as defined using DHCP option 150.
Option 66	The IP address of the TFTP server required by IP phones for configuration, as defined using DHCP option 66.
Add button	Opens the IP Pool Dialog Box, page 63-94. From here you can define a DHCP IP address pool.
Edit button	Opens the IP Pool Dialog Box, page 63-94. From here you can edit the selected IP pool.
Delete button	Deletes the selected IP pools.
Relay parameters	
Policy	The policy that DHCP relay agents implement when they receive messages already containing relay information: • Drop—The relay agent discards messages with existing relay information if option-82 information is also present. • Keep—The relay agent retains existing relay information. • Replace—The relay agent overwrites existing information with its own relay information.
Option	When selected, enables DHCP Option 82 data insertion in message requests forwarded from the DHCP client to the server. DHCP Option 82 provides the DHCP server with both the switch and port ID of the requesting client. This option makes it possible to locate where a user is physically connected to the network and prevent spoofing. See Understanding DHCP Option 82, page 63-89. When deselected, DHCP Option 82 is disabled.
Check	When selected, DHCP Option 82 reply packets received from the DHCP server are validated. Invalid messages are dropped; valid messages are stripped of the option-82 field before being forwarded to the DHCP client. When deselected, the option-82 field is removed from the packet without being checked first for validity.

DHCP Database Dialog Box

Use the DHCP Database dialog box to define external DHCP database agents that contain the automatic bindings. Each database URL that you define must be unique.

For more information, see Understanding DHCP Database Agents, page 63-88.

Navigation Path

Go to the DHCP Policy Page, page 63-92, then click the **Add** or **Edit** button beneath the Databases table.

Related Topics

- Defining DHCP Policies, page 63-90
- DHCP on Cisco IOS Routers, page 63-87
- IP Pool Dialog Box, page 63-94

Field Reference

Table 63-39 DHCP Database Dialog Box

Element	Description
Database URL	The URL of the external DHCP database agent containing the automatic bindings. The URL can be in HTTP, FTP, TFTP, or RCP format. **Note** If you define a URL, it is not necessary to define an IP address pool. However, you may do so.
Timeout	The amount of time (in seconds) the DHCP server should wait for a response from the external DHCP database agent before aborting a database transfer. The default is 300 seconds (5 minutes). **Note** A value of 0 disables the timeout.
Write Delay	The interval (in seconds) between updates sent from the DHCP server to the external DHCP database agent. The minimum delay is 60 seconds. The default is 300 seconds (5 minutes).

IP Pool Dialog Box

Use the IP Pool dialog box to define one or more address pools, which the DHCP server uses to assign dynamic addresses to DHCP clients. You must define at least one address pool, unless you have defined an external DHCP database agent.

Navigation Path

Go to the DHCP Policy Page, page 63-92, then click the **Add** or **Edit** button beneath the IP Pools table.

Related Topics

- Defining DHCP Address Pools, page 63-91
- Understanding DHCP Database Agents, page 63-88
- DHCP Database Dialog Box, page 63-94
- DHCP on Cisco IOS Routers, page 63-87

Field Reference

Table 63-40 IP Pool Dialog Box

Element	Description
Pool Name	The name of the IP pool.
Network	The IP address and subnet mask of the IP pool. This subnet contains the range of available IP addresses that the DHCP server may assign to clients. Enter an address and mask or the name of a network/host object, or click **Select** to select an object from a list or to create a new one. **Tip** You can exclude specific addresses within the range by defining them in the Excluded IPs field. See DHCP Policy Page, page 63-92.
Default Router Addresses	The IP addresses of the default routers for DHCP clients using this IP pool. After a DHCP client is booted, it begins sending packets to this router, which should be located on the same subnet as the client. Enter up to eight (8) network addresses or network/host objects, or click **Select** to select an object from a list or to create a new one.
DNS Server Addresses	The IP addresses of the DNS servers that DHCP clients using this IP pool should query when they need to correlate hostnames to IP addresses. Enter up to eight (8) network addresses or network/host objects, or click **Select** to select an object from a list or to create a new one.
NetBIOS (WINS) Server Addresses	The IP addresses of the Windows Internet Naming Service (WINS) servers used by Microsoft DHCP clients to correlate hostnames to IP addresses within a general grouping of networks. Enter up to eight (8) network addresses or network/host objects, or click **Select** to select an object from a list or to create a new one.
Domain Name	The domain name for DHCP clients using this IP pool. This name places these clients in the general grouping of networks that make up the domain.
Import All	When selected, enables remote DHCP servers to import specific DHCP options (such as the DNS server) from a centralized server. Use this option to enable configuration information to be updated automatically. When deselected, all DHCP options are local to this specific server.
Secured ARP	When selected, enables the DHCP Authorized ARP feature, which limits the leasing of IP addresses to authorized mobile users. This feature helps prevent IP spoofing by unauthorized users. See Understanding Secured ARP, page 63-89. When deselected, the DHCP Authorized ARP feature is disabled. **Note** This feature also disables dynamic ARP learning on an interface.

Table 63-40 IP Pool Dialog Box (continued)

Element	Description
Lease Never Expires	When selected, the DHCP server permanently assigns IP addresses to its clients. When deselected, addresses are leased for a predefined amount of time, as defined in the Time Length field.
Time Length (DD:HH:MM)	Applies only when the Lease Never Expires check box is deselected. The duration of the lease provided to each IP address assigned from this IP pool (using the format DD:HH:MM). After the lease expires, the assigned IP address is no longer valid and is returned to the pool.
Option 66 (IP Addresses)	The IP address of the TFTP server used to provide configuration files to IP phones. These configuration files define parameters required by IP phones to connect to Cisco CallManager. Enter up to eight (8) network addresses or network/host objects, or click **Select** to select an object from a list or to create a new one. **Note** This option is functionally similar to option 150. Either or both options may be used.
Option 150 (IP Addresses)	The IP address of the TFTP server used to provide configuration files to IP phones. These configuration files define parameters required by IP phones to connect to Cisco CallManager. Enter up to eight (8) network addresses or network/host objects, or click **Select** to select an object from a list or to create a new one. **Note** This option is functionally similar to option 66. Either or both options may be used.

NTP on Cisco IOS Routers

The Network Time Protocol (NTP) is the standard for time synchronization between network devices. Synchronized time enables you to correlate syslog and other debug output to specific events, which is essential for troubleshooting, fault analysis, and security incident tracking. Time comparisons are not possible without precise time synchronization between the logging, management, and AAA functions occurring in your network.

NTP uses the concept of a stratum to describe how far removed a machine is from an authoritative time source. For example, a stratum 1 time server is directly attached to a radio or atomic clock. NTP then distributes the time from this authoritative time source across the network. A stratum 2 time server synchronizes with a stratum 1 time server; a stratum 3 time server synchronizes with a stratum 2 time server and so on. One NTP transaction per minute is sufficient to synchronize two machines to within a millisecond.

NTP runs over the User Datagram Protocol (UDP) using port 123. Security Manager supports NTP version 3, as defined in RFC 1305.

Related Topics

- Defining NTP Servers, page 63-97

Defining NTP Servers

This procedure describes how to define the NTP servers that the routers users to synchronize time. After the NTP policy is deployed, the router uses an algorithm (based on factors such as delay, dispersion, and jitter) to determine which NTP server is the most accurate and synchronizes to that one.

At the global level, you can enable MD5 authentication and specify a source address to use on all NTP packets sent from the router.

To add an NTP server to the policy, all you need to do is enter its IP address. In addition, you can optionally define authentication parameters and determine whether a particular server should be preferred over other NTP servers of similar accuracy.

Related Topics

- NTP on Cisco IOS Routers, page 63-96

Step 1 Do one of the following:

- (Device view) Select **Platform > Device Admin > Server Access > NTP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Server Access > NTP** from the Policy Type selector. Select an existing policy or create a new one.

The NTP page is displayed. See Table 63-41 on page 63-98 for a description of the fields on this page.

Step 2 (Optional) In the Source Interface field, enter the name of the interface or interface role whose address should be used as the source interface for all NTP packets sent from the router, or click **Select** to select an interface role from a list or to create a new one. The source interface must have an IP address.

This option is useful when the NTP server cannot reach the address from which the connection originated (for example, due to a firewall). If you do not enter a value in this field, the address of the outgoing interface is used.

Note You can override this global setting for individual NTP servers, as described in Step 5.

Step 3 (Optional) Select the **Enable NTP Authentication** check box to authenticate all associations between this router and the NTP servers defined in this policy.

Step 4 Click the **Add** button under the Servers table to display the NTP Server dialog box. From here you can define an NTP server.

Step 5 Define an NTP server. See Table 63-42 on page 63-100 for a description of the available fields.

Step 6 (Optional) Define authentication parameters for this NTP server.

Note If you modify the value of a previously defined authentication key, the change affects all NTP servers that share this key.

Note When you define an authentication key in Security Manager, the value 0 is automatically appended to the end of the CLI command. This value, which represents the default authentication key encryption type, can be modified using the CLI.

Step 7 Repeat Step 5 and Step 6 to define additional NTP servers.

Step 8 Click **OK** to save your definitions locally on the client and close the dialog box. Your definitions are displayed in the Servers table.

Note To edit an NTP server, select it from the Servers table, then click **Edit**. To remove an NTP server, select it, then click **Delete**. If the key defined on the server you delete is not defined on a different NTP server, the key is also deleted.

NTP Policy Page

Use the NTP page to define one or more NTP servers that the router can use for time synchronization. This includes enabling authentication, if required, and defining a global source interface for all traffic sent to these servers.

For more information, see Defining NTP Servers, page 63-97.

Navigation Path

- (Device view) Select **Platform > Device Admin > Server Access > NTP** from the Policy selector.
- (Policy view) Select **Router Platform > Device Admin > Server Access > NTP** from the Policy Type selector. Right-click **NTP** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- NTP on Cisco IOS Routers, page 63-96
- Understanding Interface Role Objects, page 6-73
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 63-41 NTP Page

Element	Description
Source Interface	The source address for all packets sent to an NTP server. This setting might be necessary when the NTP server cannot respond to the address from which the packet originated (for example, due to a firewall). The source interface must have an IP address. If you do not define a value in this field, the address of the outgoing interface is used. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one. **Note** The source interface defined in this field is a global setting that you can override for individual NTP servers. For more information, see NTP Server Dialog Box, page 63-99.

Table 63-41 NTP Page (continued)

Element	Description
Enable NTP Authentication	When selected, enables authentication using MD5 when connecting to an NTP server. When deselected, authentication is disabled.
Servers Table	
IP Address	The IP address of the NTP server.
Source Interface	The source address for all packets sent to this NTP server. This setting overrides the global setting defined at the top of the page.
Preferred	Indicates whether this NTP server is preferred over other NTP servers of similar accuracy. **Note** By default, preferred servers are listed first in the table.
Key Number	The ID number of the key used for authentication with this NTP server.
Trusted	Indicates whether the authentication key defined for this NTP server is a trusted key.
Add button	Opens the NTP Server Dialog Box, page 63-99. From here you can define an NTP server.
Edit button	Opens the NTP Server Dialog Box, page 63-99. From here you can edit the selected NTP server.
Delete button	Deletes the selected NTP server from the table. If the key defined on the server you delete is not defined on a different NTP server, the key is also deleted.

NTP Server Dialog Box

Use the NTP Server dialog box to define the address of an NTP server that the router can use to perform time synchronization. In addition, you can use this dialog box to define a default source interface for NTP packets sent to this server and authentication parameters.

Navigation Path

Go to the NTP Policy Page, page 63-98, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining NTP Servers, page 63-97
- NTP on Cisco IOS Routers, page 63-96
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 63-42 NTP Server Dialog Box

Element	Description
IP Address	The IP address of the NTP server. Enter an address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one.
Source Interface	The source address for all packets sent to this NTP server. This setting might be necessary when the NTP server cannot respond to the address from which the packet originated (for example, due to a firewall). The source interface must have an IP address. If you do not define a value in this field and there is no global setting, the address of the outgoing interface is used. **Note** This setting overrides the global setting you defined on the NTP Policy Page, page 63-98. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one.
Preferred	When selected, this NTP server is preferred over other NTP servers of similar accuracy. If this server is used for synchronization, the time offset used to correct the local clock is calculated from this server only. **Note** If a different NTP server is significantly more accurate than the preferred server (for example, stratum 2 versus stratum 3), the router synchronizes to the more accurate server. When deselected, this NTP server is not given preference over other NTP servers of similar accuracy. The time offset used to correct the local clock is calculated by taking the combined offset of all NTP servers. We recommend that you configure an NTP server as preferred only when multiple servers have the same stratum and you can rely on the accuracy of the preferred server.
Authentication Key	The MD5 key that is used to authenticate associations with the NTP server. • Key Number—The ID number of the authentication key. Enter the key number or select a previously defined number from the list. • Key Value—An arbitrary string of up to 32 characters that defines the authentication key. Enter the string again in the Confirm field. • Trusted—When selected, this key authenticates the identity of systems attempting to synchronize with this server. When deselected, this key is not used for authentication. If you select a key number from the list and then change the key value, you are warned that saving this change affects any other NTP servers using the same authentication key. **Note** To use authentication, you must enable it from the NTP Policy Page, page 63-98.

Configuring Identity Policies

This chapter contains the following topics:

- 802.1x on Cisco IOS Routers, page 64-1
- 802.1x Policy Page, page 64-5
- Network Admission Control on Cisco IOS Routers, page 64-8
- Network Admission Control Policy Page, page 64-14

802.1x on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

The IEEE 802.1x standard defines 802.1x port-based authentication as a client-server based access control and authentication protocol that restricts unauthorized clients from connecting to a LAN through public ports. The authentication server validates each client connected to an interface before making available any services offered by the router or the LAN.

Until the client is authenticated, 802.1x access control allows only Extensible Authentication Protocol over LAN (EAPOL) traffic through the interface to which the client is connected. If authentication is successful, normal traffic can pass through the interface.

802.1x authentication provides VPN access control, enabling unauthenticated traffic to access the Internet while preventing it from accessing the VPN tunnel. This solution is especially useful for enterprises whose workers access the corporate VPN through a home access router that other family members use to access the Internet. When you use 802.1x, you create a virtual interface to carry unauthenticated traffic while authenticated traffic continues to pass through the physical interface.

802.1x requires that you use DHCP to provide IP addresses to the clients that request authentication. We recommend that you use two IP address pools, one for authenticated traffic and the other for unauthenticated traffic. If you use two pools, the DNS server in the corporate DHCP pool should point to the corporate DNS server. The DNS server for the noncorporate DHCP pool should use the DNS server provided by the ISP on the public interface. You configure DHCP by selecting a DHCP policy. See DHCP on Cisco IOS Routers, page 63-87 for more information.

Note 802.1x is supported on the following platforms—Cisco 800, 1700, 1800, 1900, 2600, 2800, 2900, 3600, 3700, 3800, 3900 Series Routers.

For more information about 802.1x, see:

- Understanding 802.1x Device Roles, page 64-2
- 802.1x Interface Authorization States, page 64-2
- Topologies Supported by 802.1x, page 64-3
- Defining 802.1x Policies, page 64-4

Understanding 802.1x Device Roles

802.1x port-based authentication uses the following device roles:

- Client—The workstation requesting access to the VPN. It must be running 802.1x-compliant client software, such as that offered with the Microsoft Windows XP operating system.
- Authentication server—Authenticates clients. The authentication server validates the client's identity and notifies the router whether the client is authorized to access the network. The Remote Authentication Dial-In User Service (RADIUS) security system with EAP extensions is the only supported authentication server. In Security Manager, a AAA (authentication, authorization, and accounting) server, as defined in a AAA server object, is the authentication server for 802.1x policies.
- Router (edge router or wireless access point)—Controls physical access to the network based on the authentication status of the client. The router is an intermediary (proxy) between the client and the authentication server, requesting identity information from the client, verifying that information with the authentication server, and relaying a response to the client. In Security Manager, the router on which you configure an 802.1x policy acts as the switch.

Related Topics

- 802.1x Interface Authorization States, page 64-2
- Topologies Supported by 802.1x, page 64-3
- Defining 802.1x Policies, page 64-4
- 802.1x on Cisco IOS Routers, page 64-1

802.1x Interface Authorization States

When you use 802.1x, the interface state determines whether to grant the client network access. By default, the interface starts in the unauthorized state. While in this state, the interface disallows all traffic in both directions, except for EAPOL packets. After a client is authenticated, the interface transitions to the authorized state, enabling all client traffic to flow normally.

If a client that does not support 802.1x is connected to an unauthorized 802.1x interface, the router requests the client's identity. In this situation, the client does not respond to the request, the interface remains in the unauthorized state, and the client is not granted access to the network. In contrast, when an 802.1x-enabled client connects to an interface that is not running the 802.1x protocol, the client initiates the authentication process by sending the EAPOL-Start frame. If no response is received, the client sends the request a fixed number of times. Because no response is received, the client begins sending frames as if the interface were in the authorized state.

You can control the interface authorization state by selecting one of the following options:

- Auto—Enables 802.1x authentication, which causes the interface to start in the unauthorized state. Only EAPOL frames are sent and received through the interface. Authentication begins when the link state of the interface transitions from down to up or when an EAPOL-Start frame is received. The router requests the client's identity and begins relaying authentication messages between the client and the authentication server. The router uses the MAC address of each client trying to access the network as unique client identifiers.
- Force authorized—Disables 802.1x authentication, which causes the interface to move to the authorized state without authenticating the client.

After a client is successfully authenticated, the interface state changes to authorized, which enables all frames from the client to enter the network. If authentication fails, the interface remains in the unauthorized state, but authentication can be retried. If the authentication server cannot be reached, the router can retransmit the request. If the authentication server does not respond after the defined number of attempts, authentication fails and network access is denied to the client.

When a client logs off, it sends an EAPOL-Logoff message, which causes the interface to return to the unauthorized state.

Related Topics

- Understanding 802.1x Device Roles, page 64-2
- Topologies Supported by 802.1x, page 64-3
- Defining 802.1x Policies, page 64-4
- 802.1x on Cisco IOS Routers, page 64-1

Topologies Supported by 802.1x

802.1x port-based authentication supports two topologies:

- Point-to-point
- Wireless LAN

In a point-to-point configuration, only one client can be connected to the 802.1x-enabled interface. The router detects the client when the interface state changes from down to up. If a client leaves the network or is replaced by another client, the interface state changes from up to down, which returns the interface to the unauthorized state.

Figure 64-1 802.1x Topology

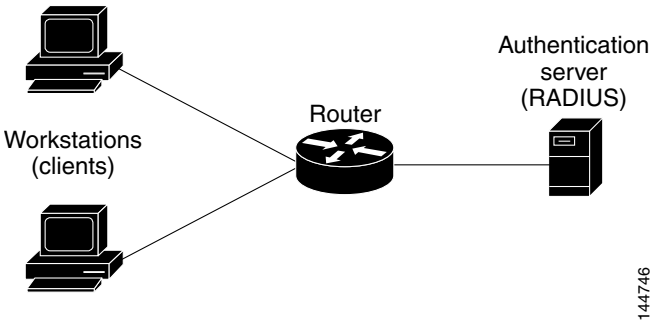

In a wireless LAN configuration, the 802.1x interface is configured in multihost mode, which is authorized as soon as one client is authenticated. After the interface is authorized, all other clients indirectly attached to the interface are granted access to the network. If the port becomes unauthorized (either because reauthentication fails or an EAPOL-Logoff message is received), the router denies access to the network to all attached clients. In this topology, the wireless access point is a client to the router and is responsible for authenticating the clients attached to it.

Related Topics

- Understanding 802.1x Device Roles, page 64-2
- 802.1x Interface Authorization States, page 64-2
- Defining 802.1x Policies, page 64-4
- 802.1x on Cisco IOS Routers, page 64-1

Defining 802.1x Policies

You configure an 802.1x policy by defining:

- The AAA server group containing the AAA server that authenticates hosts that are trying to connect to the network.
- The virtual interface that carries unauthenticated traffic and the physical interface that carries authenticated traffic.
- (Optional) Properties of the physical interface, including the control type, automatic reauthentication, and several timeout values.

If the router on which you are defining the 802.1x policy is not part of a VPN (for example, if it is directly connected to the corporate network to which you want to restrict access), you must manually define an access list. You can do this by defining an access rules policy (see Chapter 16, "Managing Firewall Access Rules").

Before You Begin

- Configure the selected router with a DHCP policy that contains two IP address pools, one for authenticated clients and one for unauthenticated clients. See Defining DHCP Policies, page 63-90.
- Make sure the router can route packets to the configured AAA (RADIUS) server. You can verify this by pinging the server from the router.

Related Topics

- Understanding 802.1x Device Roles, page 64-2
- 802.1x Interface Authorization States, page 64-2
- Topologies Supported by 802.1x, page 64-3
- 802.1x on Cisco IOS Routers, page 64-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Identity > 802.1x** from the Policy selector.
- (Policy view) Select **Router Platform > Identity > 802.1x** from the Policy Type selector. Select an existing policy or create a new one.

The 802.1x page is displayed. See Table 64-1 on page 64-6 for a description of the fields on this page.

Step 2 Enter the name of the AAA server group containing the AAA server to use for authenticating clients using 802.1x, or click **Select** to select a server group from a list or to create a new one. The selected AAA server must use RADIUS with EAP extensions.

Note Each AAA server in the selected group must be configured to communicate with an interface that exists on the router; otherwise, validation fails.

Step 3 In the Virtual Template field, enter the name of the interface or interface role that serves as the untrusted, virtual interface for carrying unauthenticated traffic, or click **Select** to select an interface role from a list or to create a new role. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Note Integrated Services Routers (ISRs), such as the Cisco 800, 1800, 1900, 2800, 2900, 3800, and 3900 Series, automatically use VLANs to carry unauthenticated traffic. If you define a virtual template, however, it is used in place of the VLAN.

Note Deployment might fail if PPP is defined on the virtual template defined here. See Defining PPP Connections, page 62-72.

Step 4 Enter the name of the interface or interface role that serves as the trusted, physical interface for carrying authenticated traffic, or click **Select** to select a role from a list.

The interface role you select should represent the internal protected interface that was configured as part of the VPN topology and no other physical interface on the selected router. For more information, see Defining the Endpoints and Protected Networks, page 25-34.

Step 5 (Optional) Modify the defaults of the physical interface used for 802.1x authentication. See Table 64-1 on page 64-6 for details.

802.1x Policy Page

Use the 802.1x policy page to create policies that limit VPN access to authorized users. Authenticated traffic is allowed to pass through a designated physical interface on the router. Unauthenticated traffic is allowed to pass through a virtual interface to the Internet but is not allowed to access the VPN.

For more information, see Defining 802.1x Policies, page 64-4.

Note 802.1x policies require DHCP address pools in order to assign IP addresses to clients. You define these pools by defining a DHCP policy on the same router. See DHCP Policy Page, page 63-92.

Navigation Path

- (Device view) Select **Platform > Identity > 802.1x** from the Policy selector.
- (Policy view) Select **Router Platform > Identity > 802.1x** from the Policy Type selector. Right-click **802.1x** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- 802.1x on Cisco IOS Routers, page 64-1
- Understanding AAA Server and Server Group Objects, page 6-27
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 64-1 802.1x Page

Element	Description
AAA Server Group	The RADIUS AAA server group that authenticates the credentials of users trying to access a VPN tunnel. Enter the name of a AAA server group object, or click **Add** to select one from a list or to create a new AAA server group object. **Note** Each AAA server in the selected group must be configured to communicate with an interface that exists on the router; otherwise, validation fails.
Virtual Template	Mandatory for all routers except Integrated Services Routers (ISRs). The untrusted, virtual interface that provides Internet access to unauthenticated traffic. Enter the name of an interface or interface role, or click **Select** to select one from a list or to create a new group object. **Note** You do not need to configure a virtual template for ISRs, because they automatically use VLANs to provide access. If you do define a virtual template, however, it is used instead of the VLAN. **Note** Deployment might fail if PPP is defined on the virtual template defined here. See PPP Dialog Box, page 62-77.
Interface	The trusted, physical interface that provides VPN access to authenticated traffic. Enter the name of an interface or interface role, or click **Select** to select one from a list or to create a new group object. If you use an interface role, the pattern defined in the interface role must represent only one physical interface on the selected device. This interface should be the internal protected interface that you configured as part of the VPN topology. For more information, see Defining the Endpoints and Protected Networks, page 25-34.
Number of retries	The number of times the physical interface resends an Extensible Authentication Protocol (EAP) request/identity frame to a client if a response is not received before restarting authentication. Valid values range from 1 to 10. The default is 2. **Note** You should change the default only to adjust for unusual circumstances, such as unreliable links or specific problems with certain clients and authentication servers.

Table 64-1 802.1x Page (continued)

Element	Description
Control type	The control state of the interface, which determines whether the host is granted access to the network. Options are: • Force Authorize—Disables 802.1x authentication and causes the interface to move to the authorized state without requiring any authentication exchange. This means the interface transmits and receives normal traffic without 802.1x-based authentication of the host. This is the default. • Auto—Enables 802.1x authentication and causes the interface to begin in the unauthorized state, allowing only EAPOL frames to be sent and received through the interface. If a host is successfully authenticated, the interface state changes to authorized, which enables all frames from the host through the interface.
Enable client reauthentication	When selected, enables periodic reauthentication of client PCs on the 802.1x interface. Reauthentication is performed after the interval defined in the Client reauthentication period timeout field. The default period is 3600 seconds (1 hour). When deselected, periodic reauthentication is not performed.
Client reauthentication period timeout	Applies only when the Enable client reauthentication check box is selected. The number of seconds between client reauthentication attempts. Valid values range from 1 to 65535 seconds. The default is 3600 seconds (1 hour).
Quiet period	The amount of time the router remains in a quiet state after a failed authentication exchange with the client. Authentication exchanges might fail, for example, because the client provided an invalid password. Valid values range from 1 to 65535 seconds. The default is 120 seconds. **Note** Entering a value smaller than the default provides a faster response time to the user.
Rate Limit period	The interval after which the interface throttles the EAP-Start packets it receives from malfunctioning client PCs. Use this setting, called rate limiting, to prevent these clients from wasting router processing power. Valid values range from 1 to 65535 seconds. By default, rate limiting is disabled. **Note** To disable an existing rate limit, delete the value defined in this field and leave the field blank.
AAA Server timeout	The number of seconds the router waits before retransmitting packets to the AAA server. If the router sends an 802.1x packet to the AAA server and the server does not respond, the router sends another packet after this interval elapses. Valid values range from 1 to 65535 seconds. The default is 30 seconds.

Table 64-1 802.1x Page (continued)

Element	Description
Supplicant period	The number of seconds the router waits before retransmitting EAP-Request/Identity packets to the supplicant (client PC). If the router sends an EAP-Request/Identity packet to the client PC (supplicant) and the supplicant does not respond, the router sends the packet again after this interval elapses. Valid values range from 1 to 65535 seconds. The default is 30 seconds.

Network Admission Control on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Network Admission Control (NAC), an industry initiative sponsored by Cisco Systems, uses the network infrastructure to enforce security-policy compliance on all devices seeking to access network computing resources, thereby limiting damage from viruses and worms. By using NAC, organizations can provide network access to endpoint devices such as PCs, PDAs, and servers that are verified to be fully compliant with established security policy. NAC can also identify noncompliant devices and deny them access, place them in a quarantined area, or give them restricted access to computing resources.

Network access decisions are made through a process of posture validation, which evaluates the posture credentials presented by the endpoint device. These credentials can include such information as the endpoint's antivirus state, operating system version, operating system patch level, or Cisco Security Agent version and settings.

You can use NAC to enforce security policy compliance in many types of deployments, including branch offices, remote access, and dial-in access.

NAC policies in Security Manager enable a Cisco IOS router to act as a Network Access Device (NAD) for enforcing policy compliance on devices seeking to access the network. The following topics describe additional details about NAC:

- Understanding NAC Components, page 64-9
- Understanding NAC System Flow, page 64-9

The following topics describe the tasks you perform to create NAC policies on Cisco IOS routers:

- Defining NAC Setup Parameters, page 64-10
- Defining NAC Interface Parameters, page 64-11
- Defining NAC Identity Parameters, page 64-13

Router Platforms Supporting NAC

To configure NAC policies on a router, the router must be running Cisco IOS Software Release 12.3(8)T images and higher (with the Advanced Security feature set). However, the following routers do not support NAC:

- Cisco 7600 Series (7603, 7604, 7606, 7609, 7613)

- Cisco 7300 Series (7301, 7304)
- Cisco 7100 Series VPN Routers (7120, 7140, 7160)
- Cisco 3600 Series Multiservice Platforms (3620, 3631, 3661, 3662)
- Cisco 1700 Series Modular Access Routers (1710, 1720, 1750)
- Cisco 1600 Series (1601, 1602, 1603, 1604, 1605)
- Cisco ASR 1000 Series Aggregation Services Routers (all models)
- Cisco 800 Series (801, 803, 805, 811, 813, 828, 851, 857, 871, 876, 877, 878)
- Cisco SOHO 90 Series Secure Broadband Routers (91, 96, 97)
- Cisco SOHO 77 Series (71, 76, 77 ADSL, 77 H ADSL, 78)

Understanding NAC Components

NAC contains the following components:

- Cisco Trust Agent (CTA)—The CTA acts as the NAC client. It provides posture credentials for the endpoint device on which it is installed, including the type of operating system and the version of antivirus software installed.
- Network access device (NAD)—The NAD initiates posture validation with the CTA when its Intercept ACL is triggered. It relays posture credentials received from the CTA to a AAA server. In return, the NAD receives configuration information from the AAA server, which it enforces on the selected interface. The NAD also:
  - Periodically polls the CTA to confirm that it is communicating with the same client at this IP address.
  - Revalidates all current sessions.
  - Sends username and password information from devices lacking a CTA (clientless hosts) to the AAA server for authentication.
  - Supports an exception list of predefined actions applied to specific devices, based on the device IP address or MAC address.

When you configure NAC policies in Security Manager, you are configuring the behavior of the Cisco IOS router acting as the NAD.

- AAA server—The AAA server obtains and validates posture credentials received from the CTA and returns the access policy to be enforced on the NAD. The AAA server must be a Cisco Secure Access Control Server (ACS), running the RADIUS protocol. Existing ACS authorization support can be used to provide access to clientless hosts. Posture validation rules and the access policies resulting from those rules are configured on the ACS.

Related Topics

- Understanding NAC System Flow, page 64-9
- Network Admission Control on Cisco IOS Routers, page 64-8

Understanding NAC System Flow

As shown in Figure 64-2, the system flow for NAC is:

1. An IP packet from a connecting device triggers the Intercept ACL configured on the NAD.

2. The NAD triggers posture validation with the CTA configured on the device using the Extensible Authentication Protocol over User Datagram Protocol, otherwise known as EAP over UDP, or simply EoU.
3. The CTA sends its posture credentials to the NAD using EAP over UDP.
4. The NAD sends these posture credentials to the ACS using RADIUS.
5. The ACS performs posture validation, which determines whether to allow the device to access the network. (If necessary, the ACS requests additional posture validation from a third-party server. For example, if the CTA forwards credentials that are specific to a particular antivirus application, the ACS forwards this information via the HCAP protocol to a vendor server for validation.) If the device is a clientless host, the ACS checks the username and password it receives against its locally stored list.
6. The ACS directs the NAD to enforce the appropriate access policy on the requesting device. Access may be granted, denied, redirected, or restricted.

Figure 64-2 NAC System Flow

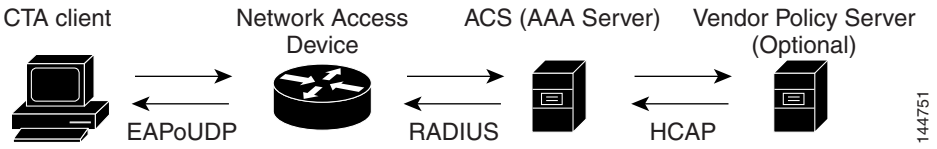

Related Topics

- Understanding NAC Components, page 64-9
- Network Admission Control on Cisco IOS Routers, page 64-8

Defining NAC Setup Parameters

You configure NAC setup parameters by selecting the AAA server groups that obtain and validate the posture credentials received from devices trying to connect to the network. You can configure an option that allows devices lacking the Cisco Trust Agent (CTA) to be authenticated by a predefined username and password stored on a Cisco Secure Access Control Server (ACS). Additionally, you can modify default settings for EAP over UDP. This is the protocol used for posture validation communications between the Cisco IOS router serving as the network access device (NAD) and the device trying to access your network.

Related Topics

- Defining NAC Interface Parameters, page 64-11
- Defining NAC Identity Parameters, page 64-13
- Network Admission Control on Cisco IOS Routers, page 64-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Identity > Network Admission Control** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Identity > Network Admission Control** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The NAC Setup tab is displayed. See Table 64-2 on page 64-15 for a description of the fields on this tab.

Step 2 Enter the name of the AAA server group containing the AAA server that performs posture validation, or click **Select** to select the server group from a list or to create a new one. The selected AAA server group must contain ACS devices running RADIUS.

Note Each AAA server in the selected group must be configured to communicate with an interface that exists on the router; otherwise, validation fails.

Step 3 (Optional) Select up to two AAA server groups as backups to the main server group. If all the servers in the main server group go down, the servers in the backup server group perform NAC.

Both backup server groups must consist of ACS devices running RADIUS.

Step 4 (Optional) Under EAP over UDP, select one or both of the following Allow parameters:

a. Select the **Allow IP Station ID** check box to include IP addresses in the RADIUS requests sent to the ACS.

b. Select the **Allow Clientless** check box to provide access to devices that do not have the CTA installed. In such cases, the ACS authenticates these devices by checking the username and password against a predefined list.

If you do not select this check box, devices without CTA are prevented from accessing the network if their traffic matches the Intercept ACL. This is because without CTA, posture validation cannot be performed.

Note This feature is not supported on routers running Cisco IOS Software Release 12.4(6)T or later.

Step 5 (Optional) Under EAP over UDP, modify the default settings related to the EAP over UDP (EoU) protocol, if required. See Table 64-2 on page 64-15 for details.

Defining NAC Interface Parameters

You configure NAC interface parameters by selecting the interfaces on which NAC is performed. You must also define the Intercept ACL, which determines which traffic on these interfaces is subject to posture validation. Additionally, you can optionally override the device-level setting for initiating EAP over UDP sessions and subject all sessions to periodic revalidation (see Defining NAC Setup Parameters, page 64-10).

A NAC policy must include at least one interface definition to function.

Before You Begin

- Select the AAA server group containing the ACS device performing posture validation. See Defining NAC Setup Parameters, page 64-10.
- Define an ACL object that defines the traffic to subject to posture validation in NAC policies. See Creating Access Control List Objects, page 6-53.
- Define an ACL object that defines the default access on the selected interface (default ACL). See Creating Access Control List Objects, page 6-53.

Related Topics

- Defining NAC Setup Parameters, page 64-10
- Defining NAC Identity Parameters, page 64-13
- Network Admission Control on Cisco IOS Routers, page 64-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Identity > Network Admission Control** from the Policy selector, then click the **Interfaces** tab in the work area.
- (Policy view) Select **Router Platform > Identity > Network Admission Control** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Interfaces** tab.

The NAC Interfaces tab is displayed. See Table 64-3 on page 64-16 for a description of the fields on this tab.

Step 2 On the NAC Interfaces tab, select an interface definition from the table, then click **Edit**, or click **Add** to create a definition. The NAC Interface Configuration dialog box appears. See Table 64-4 on page 64-17 for a description of the fields in this dialog box.

Step 3 Enter the name of the interface or interface role on which NAC is performed, or click **Select** to select an interface role from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 (Optional) Enter the name of the ACL object that acts as the intercept ACL, or click **Select** to select it from a list or to create a new object.

The intercept ACL determines which traffic on the selected interfaces is subject to posture validation before being granted access to the network. If you do not select an ACL, all traffic on the selected interfaces is subject to posture validation.

Note If you defined an authentication proxy on the same interface as a NAC interface, you must use the same intercept ACL in both policies. Otherwise, deployment might fail. For more information about authentication proxies, see Configuring AAA Rules for IOS Devices, page 15-7.

Step 5 (Optional) To override the device-level value defined for maximum attempts to initiate an EAP over UDP session, enter a new value in the EAP over UDP Max Retries field.

Step 6 (Optional) Deselect the **Enable EOU Session Revalidation** check box if you do not want the NAD to periodically revalidate all EAP over UDP sessions.

Note Subinterfaces support default values only for the options described in Step 5 and Step 6.

Step 7 Click **OK** to save your definitions locally on the client and close the dialog box. Your interface definitions appear in the table on the NAC Interfaces tab.

Defining NAC Identity Parameters

By default, any traffic over the selected interfaces that match the intercept ACL is subjected to posture validation before it is permitted to enter the network. However, you can create an exception list of predefined actions to apply to specific devices. You use identity profiles to create this exception list. Each profile contains two elements:

- A profile definition, identifies the device to which the profile applies. Devices can be identified by their IP addresses, MAC addresses, or types (for Cisco IP phones).
- An action, which defines the result when this device tries to access the network. Each action can include an ACL, a redirect URL, or both. If you do not specify an action, the default ACL is applied.

When you configure NAC identity parameters, you first define one or more identity actions and then create the identity profiles to which these actions apply. You can apply each action to multiple profiles.

Related Topics

- Defining NAC Setup Parameters, page 64-10
- Defining NAC Interface Parameters, page 64-11
- Network Admission Control on Cisco IOS Routers, page 64-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Identity > Network Admission Control** from the Policy selector, then click the **Identities** tab in the work area.
- (Policy view) Select **Router Platform > Identity > Network Admission Control** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Identities** tab.

The NAC Identities tab is displayed. See Table 64-5 on page 64-18 for a description of the fields on this tab.

Step 2 Define one or more identity actions:

a. On the NAC Identities tab, select an identity action from the lower table, then click **Add**. The NAC Identity Action dialog box appears.

b. Define an identity action. See Table 64-7 on page 64-20 for a description of the available fields.

c. Click **OK** to save your definitions and close the dialog box. The action appears in the Identity Actions table in the NAC Identities tab.

d. (Optional) Repeat a.through c. to define additional identity actions, as required.

Step 3 Define identity profiles:

a. Select an identity profile from the upper table on the NAC Identities tab, then click **Add**. The NAC Identity Profile dialog box appears. See Table 64-6 on page 64-19 for a description of the fields in this dialog box.

b. Enter the name of an identity action (as defined in Step 2) or click **Select** to display a selector.

c. Select and define a profile definition, which identifies the device to which the profile should apply.

d. Click **OK** to save your definitions and close the dialog box. The profile appears in the Identity Profiles table in the NAC Identities tab.

e. (Optional) Repeat a. through d. to define additional identity profiles, as required.

Network Admission Control Policy Page

Network Admission Control (NAC) policies enable Cisco IOS routers acting as network access devices (NADs) to enforce access privileges when an endpoint tries to connect to a network. Access decisions are made on the basis of information provided by the endpoint device, such as its current antivirus state, thus keeping insecure nodes from infecting the network.

You can configure NAC policies on a Cisco IOS router from the following tabs on the Network Admission Control policy page:

- Network Admission Control Page—Setup Tab, page 64-14
- Network Admission Control Page—Interfaces Tab, page 64-16
- Network Admission Control Page—Identities Tab, page 64-18

For more information, see Network Admission Control on Cisco IOS Routers, page 64-8.

Navigation Path

- (Device view) Select **Platform > Identity > Network Admission Control** from the Policy selector.
- (Policy view) Select **Router Platform > Identity > Network Admission Control** from the Policy Type selector. Right-click **Network Admission Control** to create a policy, or select an existing policy from the Shared Policy selector.

Network Admission Control Page—Setup Tab

Use the Network Admission Control Setup tab to select the Cisco Secure Access Control Servers used for authentication during the NAC process, as well as to define the EAP over UDP settings for communications between the NAD and the client seeking access to the network.

Navigation Path

Go to the Network Admission Control Policy Page, page 64-14, then click the **Setup** tab.

Related Topics

- Defining NAC Setup Parameters, page 64-10
- Network Admission Control Page—Interfaces Tab, page 64-16
- Network Admission Control Page—Identities Tab, page 64-18
- Understanding AAA Server and Server Group Objects, page 6-27

Field Reference

Table 64-2 Network Admission Control Setup Tab

Element	Description
AAA Server Group	The AAA server group used for NAC authentication. You must select a server group consisting of Cisco Secure Access Control Server (ACS) devices running the RADIUS protocol. Enter the name of a AAA server group object, or click **Select** to select the object from a list or to create a new one. **Note** Each AAA server in the selected group must be configured to communicate with an interface that exists on the router; otherwise, validation fails.
Backup AAA Server Group 1	The backup AAA server group in case the AAA servers in the main group are down.
Backup AAA Server Group 2	The secondary backup AAA server group in case the AAA servers in the main group and the first backup group are down.
EAP over UDP (EoU) settings	
Allow IP Station ID	When selected, enables an IP address to be included in the calling-station-id field of RADIUS requests sent to the ACS. When deselected, IP addresses are not included in the calling-station-id field of RADIUS requests sent to the ACS.
Allow Clientless	When selected, enables devices that do not have the Cisco Trust Agent (CTA) installed to be authenticated through the use of a username and password configured on the ACS. If you select this check box, enter the username and password (including confirmation) in the fields provided. When deselected, NAC prevents devices lacking the CTA from accessing the network, if their traffic matches the intercept ACL (see NAC Interface Configuration Dialog Box, page 64-17). **Note** This feature is not supported on routers running Cisco IOS Software Release 12.4(6)T or later.
Max Retry	The maximum number of retries that all NAC interfaces on this router should make when initiating an EAP over UDP session with a connecting device. Valid values range from 1 to 3. The default is 3. **Note** You can override this global value on a specific interface, if required. See Network Admission Control Page—Interfaces Tab, page 64-16.
Rate Limit	The number of EAP over UDP posture validations that the router can handle simultaneously. Additional devices cannot be validated until one or more devices drop off. Valid values range from 1 to 200. The default is 20. If you set this value to 0, rate limiting is turned off.

Table 64-2 Network Admission Control Setup Tab (continued)

Element	Description
Port	The UDP port to use for EAP over UDP sessions. Valid values range from 1 to 65535. The default is 21862. **Note** For NAC to work, the default ACL on this router must permit UDP traffic over the port designated here for EAP over UDP traffic. For more information, see Chapter 16, "Managing Firewall Access Rules".
Enable Logging	When selected, EAP over UDP events on this router are logged to the device. When deselected, EAP over UDP logging is disabled. This is the default.

Network Admission Control Page—Interfaces Tab

Use the Network Admission Control Interfaces tab to select and configure the router interfaces on which to perform NAC. This includes configuring the Intercept ACL and selected EoU interface parameters. A NAC policy must include at least one interface definition in order to function.

Navigation Path

Go to the Network Admission Control Policy Page, page 64-14, then click the **Interfaces** tab.

Related Topics

- Defining NAC Interface Parameters, page 64-11
- Network Admission Control Page—Setup Tab, page 64-14
- Network Admission Control Page—Identities Tab, page 64-18
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 64-3 Network Admission Control Interfaces Tab

Element	Description
Interfaces	The name of the interface on which NAC is being performed.
Intercept ACL	The name of the Intercept ACL, which determines the incoming traffic that triggers the interface to make a posture validation check.
EoU Max Retries	The maximum number of retries that this interface should perform when it initializes an EoU session with a connecting device.
Revalidate	Indicates whether the interface revalidates its EoU sessions to make sure they are still active.
Add button	Opens the NAC Interface Configuration Dialog Box, page 64-17. From here you can define a NAC interface.

Table 64-3 Network Admission Control Interfaces Tab (continued)

Element	Description
Edit button	Opens the NAC Interface Configuration Dialog Box, page 64-17. From here you can edit the selected NAC interface.
Delete button	Deletes the selected NAC interfaces from the table.

NAC Interface Configuration Dialog Box

Use the NAC Interface Configuration dialog box to add or edit the router interfaces on which NAC is being performed.

Navigation Path

Go to the Network Admission Control Page—Interfaces Tab, page 64-16, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining NAC Interface Parameters, page 64-11
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Creating Interface Role Objects, page 6-74
- Creating Access Control List Objects, page 6-53

Field Reference

Table 64-4 NAC Interface Configuration Dialog Box

Element	Description
Interface	The interface that will perform NAC on connecting devices. Enter the name of an interface or interface role, or click **Select** to select an object from a list or to create a new one.
Intercept ACL	The ACL that defines the traffic requiring posture validation. Enter the name of an ACL object, or click **Add** to select an object from a list or to create a new one. **Note** If an authentication proxy is configured on the same interface as NAC, the same Intercept ACL must be used in both policies. Otherwise, deployment may fail. For more information about authentication proxies, see Configuring AAA Rules for IOS Devices, page 15-7.
EAP over UDP Max Retries	The maximum number of times that the router should try to initiate an EoU session with a connecting device. Valid values range from 1 to 3. The default is 3. **Note** Subinterfaces support the default value only.
Enable EoU Session Revalidation	When selected, the router revalidates its EoU sessions as required. This is the default. When deselected, EoU session revalidation is not performed. **Note** Subinterfaces support the default value only.

Network Admission Control Page—Identities Tab

Use the Network Admission Control Identities tab to view, create, edit, and delete NAC identity profiles and identity actions. Identity profiles define a specific action to perform on traffic received from selected devices, as identified by their IP address, MAC address, or device type. In this way, devices with identity profiles are handled by NAC without having to undergo posture validation against an ACS.

Navigation Path

Go to the Network Admission Control Policy Page, page 64-14, then click the **Interfaces** tab.

Related Topics

- Defining NAC Identity Parameters, page 64-13
- Network Admission Control Page—Setup Tab, page 64-14
- Network Admission Control Page—Interfaces Tab, page 64-16
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 64-5 Network Admission Control Identities Tab

Element	Description
Identity Profiles Table	
Profile Definition	The type of identity profile—device IP address, MAC address, or device type (IP phone).
Action Name	The name of the action (defined in the Identity Actions table) that is assigned to this NAC identity profile.
Add button	Opens the NAC Identity Profile Dialog Box, page 64-19. From here you can define an identity profile.
Edit button	Opens the NAC Identity Profile Dialog Box, page 64-19. From here you can edit a selected identity profile.
Delete button	Deletes the selected identity profiles from the table.
Identity Actions Table	
Action Name	The name of the identity action.
ACL	The ACL applied to profiles to which this identity action is assigned.
Redirect URL	The URL to which traffic from devices to which this identity action is assigned are redirected.
Add button	Opens the NAC Identity Action Dialog Box, page 64-19 for defining a NAC identity action.
Edit button	Opens the NAC Identity Action Dialog Box, page 64-19 for editing a selected NAC identity action.
Delete button	Deletes the selected identity actions from the table.

NAC Identity Profile Dialog Box

Use the NAC Identity Profile dialog box to add or edit the NAC profiles assigned to devices that match a specific identity. Identity profiles define a NAC action to apply to all traffic coming from a specific device, based on its IP address, MAC address, or device type (for IP phones).

Navigation Path

Go to the Network Admission Control Page—Identities Tab, page 64-18, then click the **Add** or **Edit** button beneath the Identity Profiles table.

Related Topics

- NAC Identity Action Dialog Box, page 64-19
- Defining NAC Identity Parameters, page 64-13

Field Reference

Table 64-6 NAC Identity Profile Dialog Box

Element	Description
Action Name	The name of the action to assign to the profile. Enter the name of an action, or click **Select** to display a selector. For more information about creating actions, see NAC Identity Action Dialog Box, page 64-19.
Profile Definition	The device to which this profile is assigned: • IP Address—The IP address of the device to which this profile should be assigned. The same IP address cannot be used in more than one profile. • MAC Address—The MAC address of the device to which this profile should be assigned. • Cisco IP Phone—Used when defining a NAC identity profile for Cisco IP phones.

NAC Identity Action Dialog Box

Use the NAC Identity Action dialog box to add or edit the actions assigned to NAC identity profiles.

Navigation Path

Go to the Network Admission Control Page—Identities Tab, page 64-18, then click the **Add** or **Edit** button beneath the Identity Actions table.

Related Topics

- NAC Identity Profile Dialog Box, page 64-19
- Defining NAC Identity Parameters, page 64-13
- Creating Access Control List Objects, page 6-53

Field Reference

Table 64-7 NAC Identity Action Dialog Box

Element	Description
Name	A descriptive name for the identity action. Use this name when you select an action to assign to a NAC identity profile. See NAC Identity Profile Dialog Box, page 64-19.
Access Control Lists	The ACL that defines how to handle traffic received from a device which is assigned a profile that includes this action. Enter the name of an ACL object, or click **Add** to select an object from a list or to create a new one. **Note** You cannot select the same ACL object that is being used for the intercept ACL. See NAC Interface Configuration Dialog Box, page 64-17.
Redirect URL	The address of the remediation server to which traffic from the device should be redirected. Redirect URLs are usually of the form **http://URL** or **https://URL**.

Configuring Logging Policies

This chapter contains the following topics:

- Logging on Cisco IOS Routers, page 65-1
- Syslog Logging Setup Policy Page, page 65-7
- Syslog Servers Policy Page, page 65-10
- NetFlow Policy Page, page 65-12

Logging on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Security Manager provides the following policies for configuring logging on a Cisco IOS router:

- Syslog Logging Setup—Enable the syslog-logging feature, and define basic logging parameters. For more information, see Defining Syslog Logging Setup Parameters, page 65-1.
- Syslog Servers—Define the remote servers to which syslog messages are sent. For more information, see Defining Syslog Servers, page 65-3.
- NetFlow—Enable NetFlow logging by providing parameters and interfaces. See Defining NetFlow Parameters, page 65-6 for more information.

Note We strongly recommend configuring a Network Time Protocol (NTP) policy on all routers on which logging is enabled. NTP synchronization provides accurate timestamps for syslog messages, which is essential for comparing logs on multiple devices.

Defining Syslog Logging Setup Parameters

This procedure describes enabling syslog logging on the router, and defining which messages are sent to a syslog server. In addition, you can optionally define:

- The source interface for all syslog messages sent from this device.
- The messages that are saved to a local buffer.

- An origin identifier added to each message.
- A rate limit on the number of messages that can be sent.

Note To send syslog messages from the router to a syslog server, you must also define the IP address of the syslog server. For more information, see Defining Syslog Servers, page 65-3.

Related Topics

- Defining Syslog Servers, page 65-3
- Understanding Log Message Severity Levels, page 65-4
- Logging on Cisco IOS Routers, page 65-1

Step 1 Do one of the following to access the router's Syslog Logging Setup page:

- (Device view) Select **Platform > Logging > Syslog Logging Setup** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > Syslog Logging Setup** from the Policy Type selector. Select an existing policy or create a new one.

The Syslog Logging Setup page is displayed. See Table 65-2 on page 65-8 for a description of the fields on this page.

Step 2 Select **Enable Logging** to turn on the syslog logging feature. If this option is not selected, no log messages are created.

Tip To use the device's default logging settings, or to restore the default settings, simply select Enable Logging, ensure all other fields are blank, then click **Save**. The default settings vary by device. See your router documentation for more details.

Step 3 (Optional) In the Source Interface field, enter the name of the interface or interface role whose address should be used as the source interface for all log messages sent to a syslog server; or click **Select** to select an interface role from a list or to create a new one. The source interface must have an IP address.

This option is useful when the syslog server cannot reach the address from which the connection originated (for example, due to a firewall). If you do not enter a value in this field, the address of the outgoing interface is used.

Step 4 (Optional) To send log messages to a syslog server:

a. Select **Enable Trap**. This option is selected by default.

b. Select a value from the Trap Level list. All messages of this severity or greater (that is, having the same or a lower severity-level number) are sent to the syslog server; messages of a lesser severity are ignored. For more information about severity levels, see Table 65-1 on page 65-4.

Step 5 (Optional) To save log messages locally to a buffer on the router:

a. Select **Enable Buffer**. This option is selected by default.

b. Enter the Buffer Size in bytes.

c. Select the lowest severity level for messages to be saved to the buffer. All messages of that severity level or greater are saved to the buffer.

d. Select **Use XML Format** to save messages in XML format. (You can configure both the regular buffer and the XML buffer in the same policy.) If you select this option, enter the size of the XML buffer in bytes.

Note Make sure not to make buffers so large that the router runs out of memory for other tasks. If this happens, deployment may fail.

Step 6 (Optional) Define a rate limit to prevent a flood of output messages:

a. Select **Enable Rate Limit**. This option is selected by default.

b. Enter the maximum number of messages that can be sent per second.

c. Select the severity levels to *exclude* from the rate limit. For example, if you select 2 (critical), all syslog messages of severity levels 0-2 are sent to the syslog server regardless of the defined rate limit.

d. Select **All Messages** to apply the rate limit to all syslog messages *except* console messages (and excepting those severity levels specifically excluded above).

e. Select **Console Messages** to apply the rate limit to console messages only.

Note If you enable rate limiting without specifying any options, the default settings (10 messages per second, applied to console messages only) are applied.

Step 7 (Optional) To add an origin identifier to the beginning of each syslog message:

a. Select the type of origin ID to send—the IP address of the router, its host name, or a text string that you provide.

b. If you select String, enter the desired text in the field provided. Spaces are permitted.

The origin identifier is useful for identifying the source of syslog messages in cases where you send output from multiple devices to a single syslog server.

Note The origin identifier is not added to messages sent to local destinations, such as the buffer, the console, and the monitor.

Defining Syslog Servers

This procedure describes how to define the servers to which the router should send syslog messages. When you define a syslog server, you can choose whether the logging messages it receives should be forwarded as plain text or in XML format.

If you define multiple syslog servers, logging messages are sent to all of them.

Before You Begin

- Enable syslog logging and define basic logging parameters on the Syslog Logging Setup page. For more information, see Defining Syslog Logging Setup Parameters, page 65-1.

Related Topics

- Defining Syslog Logging Setup Parameters, page 65-1
- Understanding Log Message Severity Levels, page 65-4
- Logging on Cisco IOS Routers, page 65-1

Step 1 Do one of the following to access the router's Syslog Servers page:

- (Device view) Select **Platform > Logging > Syslog Servers** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > Syslog Servers** from the Policy Type selector. Select an existing policy or create a new one.

The Syslog Servers page is displayed. See Table 65-3 on page 65-11 for a description of the fields on this page.

Step 2 To define a server to receive syslog messages from this router, click the **Add** button below the table to open the Syslog Server dialog box. See Table 65-4 on page 65-12 for more about this dialog box.

Step 3 In the IP Address field, enter the address of the desired syslog server, or click **Select** to select a network/host object from a list or to create a new one. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 4 (Optional) Select **Forward Messages in XML Format** to forward received syslog messages in XML format instead of plain text.

Step 5 Click **OK** to save your definition and close the dialog box. The syslog server you defined is displayed in the table.

Note To edit a syslog server, select it from the table, then click **Edit**. To remove a syslog server, select it, then click **Delete**.

Understanding Log Message Severity Levels

Syslog messages on Cisco IOS routers are classified into eight severity levels. Each severity level is identified by a number and a corresponding name. The lower the number, the greater the severity, as shown in the following table.

Table 65-1 Syslog Message Severity Levels

Level Number	Level Name	Description
0	emergency	System unusable
1	alert	Immediate action needed
2	critical	Critical conditions
3	errors	Error conditions
4	warnings	Warning conditions
5	notifications	Normal but significant condition
6	informational	Informational messages only
7	debugging	Debug messages

Related Topics

- Defining Syslog Logging Setup Parameters, page 65-1
- Defining Syslog Servers, page 65-3

• Logging on Cisco IOS Routers, page 65-1

NetFlow on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

The ability to characterize IP traffic and understand how and where it flows is critical for network availability, performance and troubleshooting. Monitoring IP traffic flows facilitates accurate capacity planning, and ensures that network resources are used appropriately in support of organizational goals.

NetFlow is a logging feature available on IOS devices for recording, caching and transmitting IP traffic-flow information on a per-interface basis. The basic output of NetFlow is a flow record, where a "flow" is defined as a unidirectional stream of packets between a given source and destination—both defined by a network-layer IP address and transport-layer source and destination port numbers.

On the IOS device, NetFlow consists of two key components—a NetFlow cache which stores IP flow data, and the NetFlow export mechanism that transmits the NetFlow records to a collection server for data reporting. Thus, when enabled, NetFlow records and caches statistics for incoming and outgoing traffic flows, periodically transmitting these records from the device to a NetFlow collector, in the form of User Datagram Protocol (UDP) datagrams.

Several different formats for the export packet, or flow record, have evolved as NetFlow has matured, and these formats are commonly referred to as the NetFlow version. These versions are well documented, and include versions 1, 5, 7, and 9. The most commonly used format is NetFlow version 5, but version 9 is the latest format and has some advantages for extensibility, security, traffic analysis and multicasting.

Security Manager currently supports Traditional NetFlow on IOS devices. Traditional NetFlow provides a fixed flow record, even for version 9, meaning the device will use certain flags and predefined record combinations in generating the flow. The device configuration settings define export destinations, export interface, and certain version-specific transmission options.

More About Traffic Flows and NetFlow

Each packet that passes into or out of a router or switch is examined for a set of IP packet attributes. These attributes are the IP packet identity or "fingerprint," and they define whether the packet is unique, or related to other packets.

All packets with the same source/destination IP address, source/destination ports, protocol interface, and class of service are grouped into a flow and the packets and bytes are tallied. This method of flow determination (or "fingerprinting") is scalable because a large amount of network information can be condensed into a database of NetFlow information called the NetFlow cache.

In general, the NetFlow cache is constantly filling with flows, and software in the router or switch is searching the cache for flows that have terminated or expired, and these flows are exported to the NetFlow collector. (Unlike SNMP polling, NetFlow export periodically transmits information to the NetFlow collector.) The NetFlow collector has the job of assembling and organizing the exported flows to produce the real-time or historical reports used for traffic and security analysis.

NetFlow Summary

To summarize, the following steps outline NetFlow:

- NetFlow is configured on the router or switch to capture IP traffic flows

- Flow records are stored in the local NetFlow cache
- Periodically, approximately 30 to 50 flow records are bundled together and exported to a NetFlow collector server
- The collector software creates reports from the NetFlow data

Related Topics

- Logging on Cisco IOS Routers, page 65-1
- Defining NetFlow Parameters, page 65-6
- NetFlow Policy Page, page 65-12

Defining NetFlow Parameters

This procedure describes enabling NetFlow logging on the router.

Related Topics

- NetFlow on Cisco IOS Routers, page 65-5
- NetFlow Policy Page, page 65-12
- Logging on Cisco IOS Routers, page 65-1

Step 1 To access the router's NetFlow page, do one of the following:

- (Device view) Select **Platform > Logging > NetFlow** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > NetFlow** from the Policy Type selector. Select an existing policy or create a new one.

The router's NetFlow page is displayed. See NetFlow Policy Page, page 65-12 for complete descriptions of the fields on this page.

Step 2 On the **Setup** tab of the NetFlow page, specify global NetFlow parameters for the router:

- **Primary Destination** – Choose IP Address or Hostname from this list to enable NetFlow collection and to specify how the primary NetFlow collector will be defined. You can choose the blank entry to disable this option.
  - **IP Address** – Enter the IP address of the device hosting the primary NetFlow Collection Engine, and then enter the number of the **UDP Port** monitored by that flow collector (port numbers can range from 1 to 65535)
  - **Hostname** – Enter the fully qualified domain name of the device hosting the primary NetFlow Collection Engine, and then enter the number of the **UDP Port** monitored by that flow collector (port numbers can range from 1 to 65535)
- **Redundant Destination** – Choose IP Address or Hostname from this list to specify how the back-up NetFlow collector will be defined. You can choose the blank entry to disable this option.
  - **IP Address** – Enter the IP address of the device hosting the secondary NetFlow Collection Engine, and then enter the number of the **UDP Port** monitored by that flow collector (port numbers can range from 1 to 65535)
  - **Hostname** – Enter the fully qualified domain name of the device hosting the secondary NetFlow Collection Engine, and then enter the number of the **UDP Port** monitored by that flow collector (port numbers can range from 1 to 65535)

Note If you define a Primary and a Redundant Destination, flow data is transmitted to both.

- **Source Interface** – Specify the router interface through which flow data will be transmitted to the collector destination(s).
- **Version** – Define the record format to be used for flow data by choosing the appropriate NetFlow version number from this drop-down list. You can choose the blank entry to disable this option.
  - **1** – The original record format. No additional parameters are required.
  - **5** – The most widely adopted format; includes Border Gateway Protocol (BGP) autonomous system (AS) information and flow sequence numbers.

    If BGP is configured on your network, you can include either origin or peer AS information in the NetFlow records. Choose **origin-as** or **peer-as** from the AS Type drop-down list. You can choose the blank entry to disable this option.

    Check **Enable BGP Nexthop** to include BGP next hop information in the flow caches. (Note that with version 5, this information is visible in the caches, but it is not exported.)
  - **9** – The most-recent, template-based version; not yet fully supported.

    If BGP is configured on your network, you can include either origin or peer AS information in the NetFlow records. Choose **origin-as** or **peer-as** from the AS Type drop-down list. You can choose the blank entry to disable this option.

    Check **Enable BGP Nexthop** to include BGP next hop information in the flow records.

Note AS information collection is resource intensive, especially for origin-as. If you are not interested in monitoring peering arrangements, disabling AS collection may improve performance.

Step 3 On the **Interfaces** tab, define the interfaces for which traffic flows are to be reported.

- To add an interface, click the Add Row button to open the Add NetFlow Interface Settings dialog box. This dialog box is described in Adding and Editing NetFlow Interface Settings, page 65-15.
- To edit an existing interface, select the appropriate entry in the Interfaces table and then click the Edit Row button to open the Edit NetFlow Interface Settings dialog box (described in Adding and Editing NetFlow Interface Settings, page 65-15).
- To delete an existing interface, select that entry in the Interfaces table and then click the Delete Row button, and then confirm the deletion.

Note You can disable NetFlow data collection on an interface without deleting it. Refer to Adding and Editing NetFlow Interface Settings, page 65-15 for more information.

Syslog Logging Setup Policy Page

Use the Syslog Logging Setup page to enable syslog logging and define basic logging parameters on the selected Cisco IOS router.

For more information, see Defining Syslog Logging Setup Parameters, page 65-1.

Note We strongly recommend that you define an NTP policy on all routers on which logging is enabled in order to create accurate timestamps for each log message. For more information, see NTP Policy Page, page 63-98.

Note If you unassign a logging setup policy, the default logging configuration is restored on the device upon deployment.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog Logging Setup** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > Syslog Logging Setup** from the Policy Type selector. Right-click **Syslog Logging Setup** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Logging on Cisco IOS Routers, page 65-1
- Syslog Servers Policy Page, page 65-10
- NTP on Cisco IOS Routers, page 63-96
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 65-2 Syslog Logging Setup Page

Element	Description
Enable Logging	When selected, syslog logging is enabled on the device. When deselected, logging is disabled on the device. This is the default. **Tip** To use the device's default syslog logging settings, select the Enable Logging check box, then click **Save**, without entering additional values.
Source Interface	The source address for all outgoing log messages sent to a syslog server. This setting may be necessary when the syslog server cannot respond to the address from which the log message originated (for example, due to a firewall). If you do not define a value in this field, the address of the outgoing interface is used. Enter the name of an interface or interface role, or click **Select** to select an object from a list or to create a new one.

Table 65-2 Syslog Logging Setup Page (continued)

Element	Description
Trap	Defines which log messages are forwarded to a syslog server: • Enable Trap—When selected, log messages are sent to the syslog server. This is the default. When deselected, log messages are not sent. • Trap Level—The lowest severity level of messages that are logged and sent to the syslog server. All messages of this severity and greater are logged. Severity levels are identified by a name and a number. For more information, see Table 65-1 on page 65-4. **Tip** To restore the router's default trap settings, select **Enable Trap**, then select the blank setting from the Trap Level list.
Logging Buffer	Defines whether log messages are saved locally to a buffer on the device. • Enable Buffer—When selected, log messages are saved to a buffer on the device. This is the default. When deselected, a log buffer is not maintained on the device. • Buffer Size—The size of the buffer in bytes. Valid values range from 4096 to 4294967295 bytes (4 kilobytes to 4 gigabytes). The default size varies by platform. Make sure not to make the buffer so large that the router runs out of memory for other tasks; otherwise, deployment might fail. **Note** The maximum buffer size might be smaller on some devices. • Severity Level—The lowest severity level of messages that are saved in the buffer. All messages of this severity and greater are saved. On most Cisco IOS routers, the default severity level is 7 (debugging). Severity levels are identified by a name and a number. For more information, see Table 65-1 on page 65-4. • Use XML Format—When selected, log messages are saved to a buffer in XML format. (You can configure both the regular buffer and the XML buffer in the same policy.) When deselected, an XML buffer is not maintained on the device. • Buffer Size—The size of the XML buffer in bytes. Valid values range from 4096 to 4294967295 bytes (4 kilobytes to 4 gigabytes). **Note** The maximum buffer size might be smaller on some devices. **Tip** To restore the router's default buffer settings, select **Enable Trap**, erase the buffer size setting, then select the blank setting from the Severity Level list.

Table 65-2 Syslog Logging Setup Page (continued)

Element	Description
Rate Limit	Limits the rate of log messages sent to the syslog server. • Enable Rate Limit—When selected, the rate limit is enabled. When deselected, the rate limit is disabled. • Messages per Sec.—The maximum number of logging messages that can be sent per second. Valid values range from 1 to 10000. The default is 10 messages per second. • Exclude—The types of messages to *exclude* from the rate limit. This setting excludes the severity level you select as well as all messages with a lower severity level number (that is, more severe). The default is 3 (errors), which excludes all log messages with a severity level of 3, 2 (critical), 1 (alerts), or 0 (emergencies) from the rate limit. For more information about severity levels, see Table 65-1 on page 65-4. • All Messages—When selected, the rate limit applies to all messages except console messages. • Console Messages—When selected, the rate limit applies to console messages only. **Tip** To restore the router's default rate limit settings, select the **Enable Rate Limit** check box, then erase the rate limit value setting.
Origin ID	The origin identifier that is added to the beginning of all syslog messages sent from this device to the remote syslog server. The origin identifier is useful in cases where you send output from multiple devices to a single syslog server. • ID Type—The type of origin identifier added to the beginning of each syslog message. Options are: – IP Address—The IP address of the source device. – Hostname—The hostname of the source device. – String—User-defined text. • Value—Applies only when you select String as the ID type. Enter the text of the user-defined string. Spaces are permitted, except for the first character. **Note** The origin identifier is not added to messages sent to local destinations, such as the buffer, the console, and the monitor.

Syslog Servers Policy Page

Use the Syslog Servers page to create, edit, and delete servers that collect log messages from the router.

For more information, see Defining Syslog Servers, page 65-3.

Note To enable logging to the syslog servers defined on this page, you must enable logging and define basic parameters on the Syslog Logging Setup Policy Page, page 65-7.

Navigation Path

- (Device view) Select **Platform > Logging > Syslog Servers** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > Syslog Servers** from the Policy Type selector. Right-click **Syslog Servers** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Logging on Cisco IOS Routers, page 65-1
- Syslog Server Dialog Box, page 65-11
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 65-3 Syslog Servers Page

Element	Description
IP Address	The name of the syslog server, as represented by a network/host object, or its IP address.
XML	Indicates whether the syslog server receives log messages in XML format.
Add button	Opens the Syslog Server Dialog Box, page 65-11. From here you can define a syslog server.
Edit button	Opens the Syslog Server Dialog Box, page 65-11. From here you can edit the selected syslog server.
Delete button	Deletes the selected syslog server from the table.

Syslog Server Dialog Box

Use the Syslog Server dialog box to define the server that collects syslog messages from the router. You can also define whether the log messages it receives are in XML format or plain text.

Note To enable logging to the syslog servers defined on this page, you must enable logging and define basic parameters on the Syslog Logging Setup Policy Page, page 65-7.

Navigation Path

Go to the Syslog Servers Policy Page, page 65-10, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining Syslog Servers, page 65-3
- Logging on Cisco IOS Routers, page 65-1

- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 65-4 Syslog Server Dialog Box

Element	Description
IP Address	The IP address of the syslog server. Enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one.
Forward Messages in XML Format	When selected, log messages are sent to the syslog server in XML format. When deselected, log messages are sent to the syslog server as plain text.

NetFlow Policy Page

Use the NetFlow page to enable NetFlow recording and define its parameters on the selected Cisco IOS router.

The NetFlow page consists of two tabbed panels: Setup and Interfaces. The Setup tab provides global configuration parameters for NetFlow collection on the router. The Interfaces tab lists router interfaces for which NetFlow data collection is configured, and allows enabling and disabling ingress and egress accounting on a per-interface basis.

Note We strongly recommend that you define an NTP policy on all routers on which logging is enabled in order to create accurate timestamps for each log message. For more information, see NTP Policy Page, page 63-98.

Navigation Path

- (Device view) Select **Platform > Logging > NetFlow** from the Policy selector.
- (Policy view) Select **Router Platform > Logging > NetFlow** from the Policy Type selector. Select an existing policy from the Shared Policy selector, or right-click **NetFlow** to create a new policy.

Related Topics

- NetFlow on Cisco IOS Routers, page 65-5
- Defining NetFlow Parameters, page 65-6
- Adding and Editing NetFlow Interface Settings, page 65-15
- Logging on Cisco IOS Routers, page 65-1
- NTP on Cisco IOS Routers, page 63-96

Field Reference

Table 65-5 NetFlow Page

Element	Description
Setup tab	
Primary Destination Redundant Destination	The primary and secondary NetFlow collector. You must select a primary collector to enable NetFlow data collection on this device. To disable transmission of NetFlow data to either of these collectors, choose the blank entry from the drop-down list. Select whether to identify the NetFlow collector using its IP address or host name, then configure the required fields for each option: • **IP Address**—Enter the IP address of the device hosting the primary NetFlow Collection Engine. You can also specify a network/host object that specifies the IP address, or click **Select** to select the object from a list or to create a new one. In the **UDP Port** field, enter the port number monitored by the flow collector (port numbers can range from 1 to 65535). You can enter a number or the name of a port list object, or click **Select** to select an object from a list or to create a new one. • **Hostname**—Enter the fully qualified domain name of the device hosting the primary NetFlow Collection Engine. You also must specify the UDP port as you do when specifying the IP address.
Source Interface	The router interface through which flow data will be transmitted to the collector destinations. Enter an interface or interface role name, or click **Select** to select an object from a list or to create a new one.

Table 65-5 NetFlow Page (continued)

Element	Description
Version	The NetFlow version number, which defines the record format to be used for flow. You can choose the blank entry to disable this option. • **1**—The original record format. No additional parameters are required. • **5**—The most widely adopted format; includes Border Gateway Protocol (BGP) autonomous system (AS) information and flow sequence numbers. If BGP is configured on your network, you can include either origin or peer AS information in the NetFlow records. Choose **origin-as** or **peer-as** from the AS Type drop-down list. You can choose the blank entry to disable this option. Check **Enable BGP Nexthop** to include BGP next hop information in the flow caches. (Note that with version 5, this information is visible in the caches, but it is not exported.) • **9**—The most-recent, template-based version; not yet fully supported. If BGP is configured on your network, you can include either origin or peer AS information in the NetFlow records. Choose **origin-as** or **peer-as** from the AS Type drop-down list. You can choose the blank entry to disable this option. Check **Enable BGP Nexthop** to include BGP next hop information in the flow records. **Note** AS information collection is resource intensive, especially for origin-as. If you are not interested in monitoring peering arrangements, disabling AS collection might improve performance.
Interfaces tab	
Interface	The names of the interfaces on which NetFlow collection is configured.
Enable Ingress	"Enabled" indicates flow recording is enabled on this interface for incoming traffic; "Disabled" indicates incoming traffic is not recorded for this interface.
Enable Egress	"Enabled" indicates flow recording is enabled on this interface for outgoing traffic; "Disabled" indicates outgoing traffic is not recorded for this interface.
Add Row	Click this button to open the Add NetFlow Interface Settings dialog box. Adding a NetFlow interface is described in Adding and Editing NetFlow Interface Settings, page 65-15.
Edit Row	Click this button to open the Edit NetFlow Interface Settings dialog box for the selected interface. Editing NetFlow interfaces is described in Adding and Editing NetFlow Interface Settings, page 65-15.
Delete Row	Click this button to delete the selected interface. You will be asked to confirm the deletion.

Adding and Editing NetFlow Interface Settings

Use the Add NetFlow Interface Settings and Edit NetFlow Interface Settings dialog boxes to enable and disable NetFlow ingress and egress reporting for specific router interfaces.

Note Except for their titles, these two dialog boxes are identical. The following information applies to both.

Navigation Path

Go to the NetFlow Policy Page, page 65-12, then click the **Add Row** or **Edit Row** button beneath the table.

Related Topics

- Defining NetFlow Parameters, page 65-6
- Logging on Cisco IOS Routers, page 65-1

Field Reference

Table 65-6 Add/Edit NetFlow Interface Settings Dialog Box

Element	Description
Interface	The name of the interface or interface role. Enter a name or click **Select** to select an interface role from a list or to create a new one.
Enable Ingress Accounting	When this option is selected, NetFlow records are collected for traffic arriving on this interface. Deselect this option to halt data collection on this interface for incoming traffic.
Enable Egress Accounting	When this option is selected, NetFlow records are collected for traffic departing from this interface. Deselect this option to halt data collection on this interface for outgoing traffic.

Configuring Quality of Service

This chapter contains the following topics:

- Quality of Service on Cisco IOS Routers, page 66-1
- Quality of Service Policy Page, page 66-19

Quality of Service on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Quality of service (QoS) refers to the ability of a network to provide priority service to selected network traffic over various underlying technologies, including Frame Relay, ATM, Ethernet and 802.1 networks, SONET, and IP-routed networks. QoS features enhance the predictability of network service by:

- Supporting dedicated bandwidth.
- Improving loss characteristics.
- Avoiding and managing network congestion.
- Shaping network traffic.
- Setting traffic priorities across the network.

QoS is generally used at entry points to service providers, as well as at consolidation points where multiple lines converge. QoS is also useful where speed mismatches occur (for example, at the boundary between a WAN and a LAN), as these places are often traffic congestion points.

QoS policies in Security Manager are based on the Cisco Systems Modular QoS CLI (MQC). MQC standardizes the CLI and semantics for QoS features across all platforms supported by Cisco IOS software, which provides a modular and highly extensible framework for deploying QoS. Security Manager provides an easy-to-use interface for MQC that concentrates key QoS features inside a single dialog box, streamlining the creation of QoS policies for selected traffic entering and leaving the router.

For a description of the procedure for defining a QoS policy in Security Manager, see Defining QoS Policies, page 66-10.

Related Topics

- Quality of Service and CEF, page 66-2
- Understanding Marking Parameters, page 66-3

- Understanding Queuing Parameters, page 66-4
- Understanding Policing and Shaping Parameters, page 66-6

Quality of Service and CEF

Cisco Express Forwarding (CEF) is an advanced Layer 3 IP switching technology that optimizes network performance and scalability for all kinds of networks. It defines the fastest method by which a Cisco IOS router forwards packets from ingress to egress interfaces.

Certain QoS features configurable in Security Manager, such as Class-Based Policing and Class-Based Weighted Random Early Detection, are supported only on routers that run CEF. All routers from the Cisco 800 Series to the Cisco 7200 Series require CEF for these QoS features; the Cisco 7500 Series requires distributed CEF (dCEF).

Note For a complete list, see *When is CEF Required for Quality of Service* on Cisco.com at this URL: http://www.cisco.com/en/US/tech/tk39/tk824/technologies_tech_note09186a0080094978.shtml

By default, CEF is enabled as part of the router's initial configuration. To verify whether CEF is enabled on your router, use the **show ip cef** command. You can configure CEF using the CEF interface settings policy (see CEF Interface Settings on Cisco IOS Routers, page 62-25). Be aware, however, that if your router does not have CEF enabled, activating CEF could have a significant impact on your router's packet streaming. Consult your router documentation before enabling CEF.

Related Topics

- Quality of Service on Cisco IOS Routers, page 66-1

Understanding Matching Parameters

You define matching parameters by identifying the traffic on which QoS is performed, that is, classifying the interesting packets. Various classification tools are available, including protocol type, IP precedence (IPP) value, Differentiated Service Code Point (DSCP) value, and ACLs.

Traffic classes consist of a series of match criteria and a means of evaluating these criteria. For example, you might define a class with matching criteria based on several specified protocols and a DSCP value. You can then specify that a packet must match only one of these defined criteria to be considered part of this class. Your other option is to specify that packets must match all defined criteria considered part of the traffic class.

Packets that are members of a defined traffic class are forwarded according to the QoS specifications that you defined in the policy map. Packets that fail to meet any of the matching criteria are classified as members of the default traffic class.

For information about defining matching parameters in a QoS policy, see Defining QoS Class Matching Parameters, page 66-13.

Related Topics

- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Understanding Marking Parameters

Marking parameters enable you to classify packets, which entails using a traffic descriptor to categorize a packet within a specific group. This defines the packet and makes it accessible for QoS handling on the network. Both traffic policers and traffic shapers use the packet classification to ensure adherence to the contracted level of service agreed upon between the source and your network. Additionally, marking parameters enable you to take packets that might have arrived at the device with one QoS classification and reclassify them. Downstream devices use this new classification to identify the packets and apply the appropriate QoS functions to them.

Security Manager uses two types of marking for IPv4 packets—one based on IPP classes and one based on DSCP values. IPP is based on the three most significant bits in the Type of Service (ToS) byte of each packet, which means you can partition traffic into eight classes. For historical reasons, each precedence value corresponds with a name, as defined in RFC 791. Table 66-1 on page 66-3 lists the numbers and their corresponding names, from least to most important.

Table 66-1 IP Precedence Classes

Class	Name
0	routine
1	priority
2	immediate
3	flash
4	flash-override
5	critical
6	internet
7	network

Note Classes 6 and 7 are generally reserved for network control information, such as routing updates.

DSCP is based on the six most significant bits in the ToS byte (the remaining two bits are used for flow control), with values ranging from 0 to 63. The DSCP bits contains the IPP bits, which makes DSCP backward-compatible with IPP.

Marking is generally used on devices that are close to the network edge or administrative domain so that subsequent devices can provide service based on the classification mark.

For information about defining marking parameters in a QoS policy, see Defining QoS Class Marking Parameters, page 66-15.

Related Topics

- Understanding Queuing Parameters, page 66-4
- Understanding Policing and Shaping Parameters, page 66-6
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Understanding Queuing Parameters

Queuing manages congestion on traffic leaving a Cisco IOS router by determining the order in which to send packets out over an interface, based on priorities you assign to those packets. Queuing makes it possible to prioritize traffic to satisfy time-critical applications, such as desktop video conferencing, while still addressing the needs of less time-dependent applications, such as file transfer.

During periods of light traffic, that is, when no congestion exists, packets are sent out as soon as they arrive at an interface. However, during periods of transmission congestion at the outgoing interface, packets arrive faster than the interface can send them. By using congestion management features such as queuing, packets accumulating at the interface are queued until the interface is free to send them. They are then scheduled for transmission according to their assigned priority and the queuing mechanism configured for the interface. The router determines the order of packet transmission by controlling which packets are placed in which queue and how queues are serviced with respect to one another.

Security Manager uses a form of queuing called Class-Based Weighted Fair Queuing (CBWFQ). With CBWFQ, you define traffic classes based on match criteria. Packets matching the criteria constitute the traffic for this class. A queue is reserved for each class, containing the traffic belonging to that class. You assign characteristics to queues, such as the bandwidth (fixed or minimum) assigned to it and the queue limit, which is the maximum number of packets allowed to accumulate in the queue.

When you use CBWFQ, the sum of all bandwidth allocation on an interface cannot exceed 75 percent of the total available interface bandwidth. The remaining 25 percent is used for other overhead, including Layer 2 overhead, routing traffic, and best-effort traffic. Bandwidth for the CBWFQ default class, for instance, is taken from the remaining 25 percent.

For more information about queuing, see:

- Tail Drop vs. WRED, page 66-4
- Low-Latency Queuing, page 66-5
- Default Class Queuing, page 66-6

For information about defining queuing parameters in a QoS policy, see Defining QoS Class Queuing Parameters, page 66-16.

Related Topics

- Understanding Marking Parameters, page 66-3
- Understanding Policing and Shaping Parameters, page 66-6
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Tail Drop vs. WRED

After a queue reaches its configured queue limit, the arrival of additional packets causes tail drop or packet drop to take effect, depending on how you configured the QoS policy. Tail drop, which is the default response, treats all traffic equally and does not differentiate between different classes of service. When tail drop is in effect, packets are dropped from full queues until the congestion is eliminated and the queue is no longer full. This often leads to global synchronization, in which a period of congestion is followed by a period of underutilization, as multiple TCP hosts reduce their transmission rates simultaneously.

A more sophisticated approach to managing queue congestion is offered by Cisco's implementation of Random Early Detection, called Weighted Random Early Detection, or WRED. As shown in Figure 66-1, WRED reduces the chances of tail drop by selectively dropping packets when the output

interface begins to show signs of congestion. By dropping some packets early instead of waiting until the queue is full, WRED avoids dropping large numbers of packets at once and allows the transmission line to be used fully at all times.

Figure 66-1 Weighted Random Early Detection

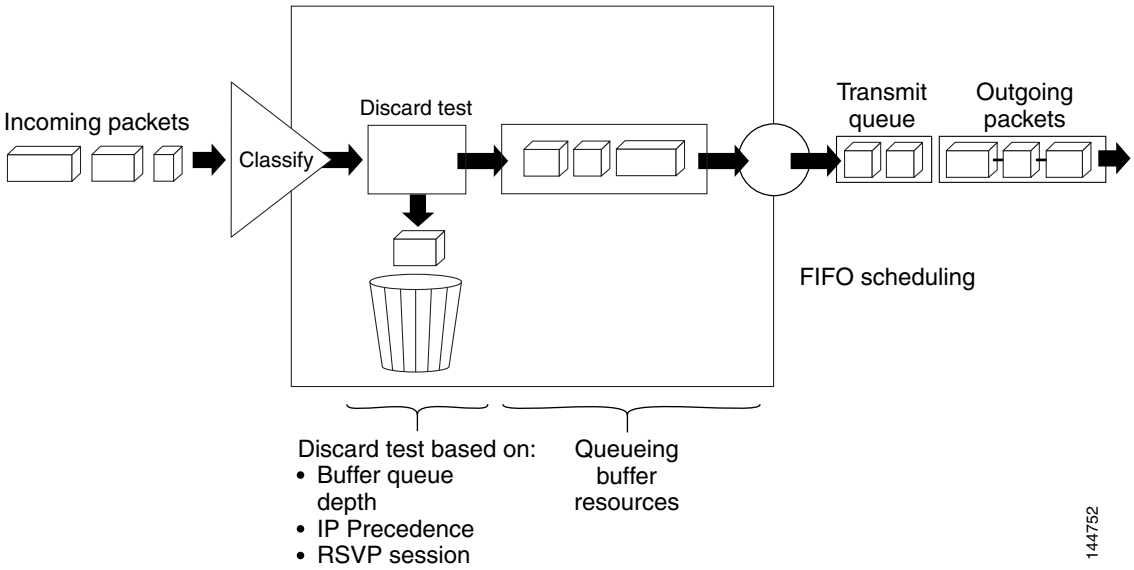

WRED is useful only when the bulk of the traffic is TCP/IP traffic, because TCP hosts reduce their transmission rate in response to congestion. With other protocols, packet sources might not respond, or might resend dropped packets at the same rate. As a result, dropping packets does not decrease congestion.

Note WRED treats non-IP traffic as precedence 0, the lowest precedence value. Therefore, non-IP traffic is more likely to be dropped than IP traffic.

Related Topics

- Low-Latency Queuing, page 66-5
- Default Class Queuing, page 66-6
- Understanding Queuing Parameters, page 66-4

Low-Latency Queuing

The low-latency queuing (LLQ) feature brings strict priority queuing to CBWFQ. Strict priority queuing gives delay-sensitive data, such as voice traffic, preference over other traffic.

Note Although it is possible to assign various types of real-time traffic to the strict priority queue, we strongly recommend that you direct only voice traffic to it.

LLQ defines the maximum bandwidth that you can allocate to priority traffic during times of congestion. Setting a maximum ensures that nonpriority traffic does not starve (meaning that this traffic is also provided with bandwidth). When the device is not congested, the priority class traffic is allowed to exceed its allocated bandwidth. Policing drops packets from the priority queue; therefore, neither WRED nor tail drop (as configured in the Queue Limit field) is used.

When LLQ is not used, CBWFQ provides weighted fair queuing based on defined classes, with no strict priority queue available for real-time traffic.

Related Topics

- Tail Drop vs. WRED, page 66-4
- Default Class Queuing, page 66-6
- Understanding Queuing Parameters, page 66-4

Default Class Queuing

You use the Fair Queue field to define the number of dynamic queues that should be reserved for the default class to use. This is the class to which traffic that does not satisfy the match criteria of other classes is directed. By default, the number of queues that are created is based on the interface bandwidth.

Table 66-2 on page 66-6 lists the default number of dynamic queues that CBWFQ uses when it is enabled on an interface:

Table 66-2 Default Number of Queues for Default Class

Bandwidth Range	Number of Dynamic Queues
Less than or equal to 64 kbps	16
More than 64 kbps and less than or equal to 128 kbps	32
More than 128 kbps and less than or equal to 256 kbps	64
More than 256 kbps and less than or equal to 512 kbps	128
More than 512 kbps	256

Related Topics

- Tail Drop vs. WRED, page 66-4
- Low-Latency Queuing, page 66-5
- Understanding Queuing Parameters, page 66-4

Understanding Policing and Shaping Parameters

Security Manager offers two kinds of traffic regulation mechanisms:

- The rate-limiting feature of Class-Based Policing for policing traffic. Policing limits traffic flow to a configured rate. Policing can be performed on a selected interface or on the control plane. See Understanding Control Plane Policing, page 66-9.

- Distributed Traffic Shaping (DTS) for shaping traffic. Traffic shaping enables you to control the traffic leaving an interface (output traffic) in order to match its flow to the speed of the remote target interface and to ensure that the traffic conforms to the policies defined for it. By shaping traffic to meet downstream requirements, you can eliminate bottlenecks in topologies with data-rate mismatches. Shaping can either be performed on selected QoS classes or at the interface level (hierarchical shaping).

Both policing and shaping mechanisms use the traffic descriptor for a packet—indicated by the classification of the packet (see Understanding Marking Parameters, page 66-3)—to ensure adherence to the agreed upon level of service. Although policers and shapers usually identify traffic descriptor violations in the same way, they differ in the way they respond to violations, as shown in Figure 66-2:

- A policer typically drops excess traffic. In other cases, it transmits the traffic with a different (usually lower) priority.
- A shaper typically delays excess traffic using a buffer, or queuing mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected.

Figure 66-2 Traffic Policing vs. Traffic Shaping

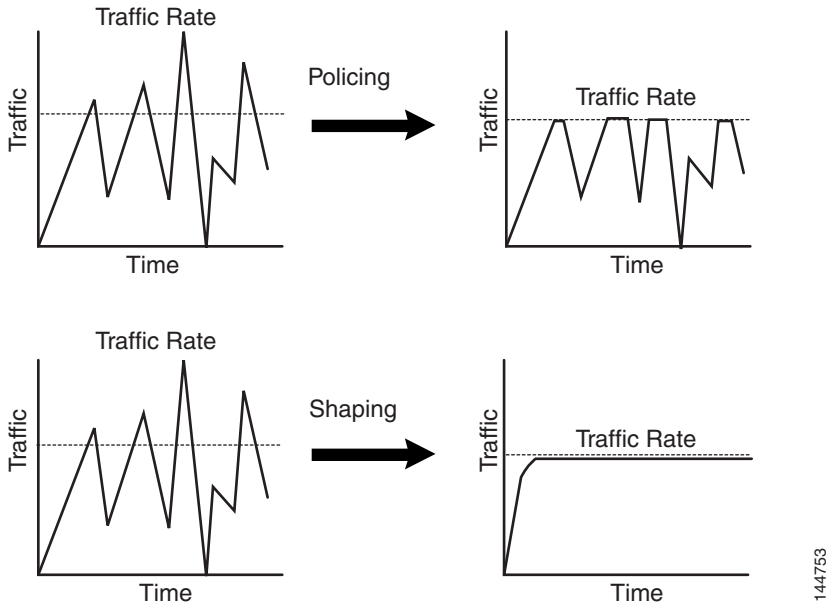

For information about defining policing and shaping parameters in a QoS policy, see Defining QoS Class Policing Parameters, page 66-17 and Defining QoS Class Shaping Parameters, page 66-18.

Related Topics

- Understanding the Token-Bucket Mechanism, page 66-8
- Understanding Marking Parameters, page 66-3
- Understanding Queuing Parameters, page 66-4
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Understanding the Token-Bucket Mechanism

Both policing and shaping use a token-bucket mechanism to regulate data flow. A token bucket is a formal definition of a rate of transfer. It has three components: a burst size, a mean rate, and a time interval (Tc). Any two values may be derived from the third using this formula:

mean rate = burst size / time interval

These terms are defined as follows:

- Mean rate—Also called the committed information rate (CIR), it specifies how much data can be sent or forwarded per unit time on average. The CIR is defined either as an absolute value or as a percentage of the available bandwidth on the interface. When defined as a percentage, the equivalent value in bits per second (bps) is calculated after deployment based on the interface bandwidth and the percent value defined in the policy.

  Note If the interface bandwidth changes (for example, more bandwidth is added), the bps value of the CIR is recalculated based on the revised amount of bandwidth.

- Burst size—Also called the committed burst (Bc) size, it specifies for each burst how much data can be sent within a given time without creating scheduling concerns. When you use percentages to calculate the CIR, burst size is measured in milliseconds.
- Time interval—Also called the measurement interval, it specifies the amount of time in seconds per burst. Over any integral multiple of this interval, the bit rate of the interface does not exceed the mean rate. The bit rate, however, might be arbitrarily fast within the interval.

In the token-bucket metaphor, tokens are put into the bucket at a certain rate. These tokens represent permission for the source to send a certain number of bits into the network. To send a packet, the regulator (policer or shaper) must remove a number of tokens from the bucket that equals the packet size.

Security Manager uses a two-bucket algorithm, as shown in Figure 66-3. The first bucket is the conform bucket and the second bucket is the exceed bucket. The full size of the conform bucket is the number of bytes specified as the normal burst size. The full size of the exceed bucket is the number of bytes specified in the maximum burst size. Both buckets are initially full, and they are updated based on the token arrival rate, which is determined by the CIR. If the number of bytes in the arriving packet is less than the number of bytes in the conform bucket, the packet conforms. The required number of tokens are removed from the conform bucket and the defined conform action is taken (for example, the packet is transmitted). The exceed bucket is unaffected.

If the conform bucket does not contain sufficient tokens, the excess token bucket is checked against the number of bytes in the packet. If enough tokens are present in the two buckets combined, the exceed action is taken on the packet and the required number of bytes are removed from each bucket. If the exceed bucket contains an insufficient number of bytes, the packet is in violation of the burst limits and the violate action is taken on the packet.

Figure 66-3 Two-Token Bucket Algorithm

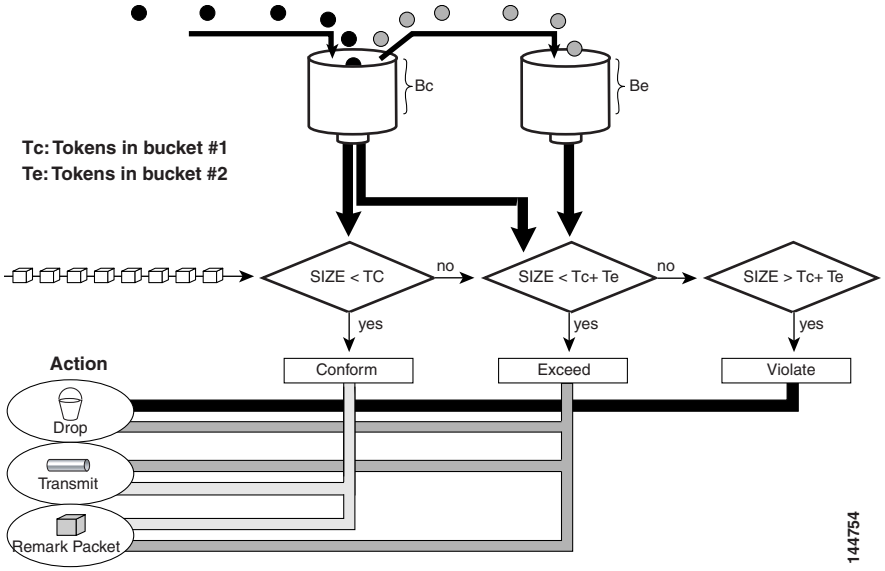

When you use traffic policing, the token-bucket algorithm provides three actions for each packet: a conform action, an exceed action, and an optional violate action. For instance, packets that conform can be configured to be transmitted, packets that exceed can be configured to be sent with a decreased priority, and packets that violate can be configured to be dropped.

Traffic policing is often configured on interfaces at the edge of a network to limit the rate of traffic entering or leaving the network. In the most common traffic policing configurations, traffic that conforms is transmitted and traffic that exceeds is sent with a decreased priority or is dropped. You can change these configuration options to suit your network needs.

When you use traffic shaping, the token-bucket mechanism includes a data buffer for holding packets that cannot be sent immediately. (Policers do not have such a buffer.) The token buckets permit packets to be sent in bursts, but places bounds on this capability so that the flow is never faster than the capacity of the buckets plus the time interval multiplied by the refill rate. The buffer also guarantees that the long-term transmission rate does not exceed the CIR.

Related Topics

- Understanding Control Plane Policing, page 66-9
- Understanding Policing and Shaping Parameters, page 66-6

Understanding Control Plane Policing

The Control Plane Policing feature enables you to manage input traffic entering the control plane (CP) of the router. The CP is a collection of processes that run at the process level on the route processor. These processes collectively provide high-level control for most Cisco IOS functions. Control plane policing protects the CP of Cisco IOS routers and switches against reconnaissance and denial-of-service (DoS) attacks, enabling the CP to maintain packet forwarding and protocol states despite an attack or heavy traffic load on the router or switch.

The Control Plane Policing feature treats the CP as a separate entity with its own ingress (input) and egress (output) ports, enabling you to use Security Manager to configure QoS policies on input. These policies are applied when a packet enters the CP. You can configure a QoS policy to prevent unwanted

packets from progressing after a specified rate limit is reached. For example, a system administrator can limit all TCP/SYN packets that are destined for the CP to a maximum rate of 1 megabit per second. Additional packets beyond this limit are silently discarded.

The following types of Layer 3 packets are forwarded to the CP and processed by aggregate control plane policing:

- Routing protocol control packets
- Packets destined for the local IP address of the router
- Packets from management protocols, such as SNMP, Telnet, and secure shell (SSH).

Note Support for output policing is available only in Cisco IOS Release 12.3(4)T and later T-train releases.

For information about how to define Control Plane Policing, see Defining QoS on the Control Plane, page 66-12. For more information about this feature, refer to the document, *Control Plane Policing* on Cisco.com at this URL:

http://www.cisco.com/en/US/docs/ios/qos/configuration/guide/ctrl_plane_policng.html

Related Topics

- Understanding the Token-Bucket Mechanism, page 66-8
- Understanding Policing and Shaping Parameters, page 66-6

Defining QoS Policies

When you define QoS policies, you must first decide whether to configure the policy on specific interfaces or on the control plane. This initial choice determines how you configure the rest of the policy, as described in the following topics:

- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12

Note If you define a QoS policy on both the interfaces and the control plane of the same device, only the control plane configuration is deployed.

Related Topics

- Quality of Service on Cisco IOS Routers, page 66-1

Defining QoS on Interfaces

You can create multiple QoS interface definitions, each of which applies to either input traffic (entering the router) or output traffic (exiting the router).

When you create a QoS interface definition on output traffic, you have the option of configuring hierarchical shaping on the interface as a whole instead of configuring shaping on individual QoS classes.

After you create your interface definitions, you must define one or more QoS classes on each interface. QoS classes contain the matching criteria that determine which packets are included in the class and the QoS functions (marking, queuing, policing, and shaping) to apply to that traffic. You can configure each interface (or interface role) with up to 16 QoS classes, each containing its own set of matching criteria and a defined set of QoS functions to apply to the traffic in that class.

For each interface, we recommend that for each interface you define at least one QoS class and a default class. If you do not configure a default class, packets that do not match the criteria of the other defined classes are treated as members of a default class that has no configured QoS functionality. Packets assigned to this class are placed in a simple first-in first-out (FIFO) queue, and are forwarded at a rate determined by the available underlying link bandwidth. This FIFO queue is managed by tail drop, which avoids congestion by dropping packets from the queue until it is no longer full.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Before You Begin

Ensure that Cisco Express Forwarding (CEF) is enabled on the router. For more information, see CEF Interface Settings on Cisco IOS Routers, page 62-25.

Related Topics

- Defining QoS Policies, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Select an existing policy or create a new one.

The Quality of Service page is displayed. See Table 66-3 on page 66-20 for a description of the fields on this page.

Step 2 In the Applied to field, select **Interfaces** to define QoS parameters for specific interfaces on the selected router.

Step 3 Click the **Add** button under the upper table to display the QoS Policy dialog box. See Table 66-4 on page 66-21 for a description of the fields in this dialog box.

Step 4 In the Interface field, enter the name of an interface or interface role, or click **Select** to display a selector.

Tip If the interface role you want is not listed in the selector, click the **Create** button or the **Edit** button to open the Interface Role Dialog Box, page 6-75. From here you can define an interface role to use in the policy.

Step 5 Select the traffic direction on which you want to apply the QoS definition, Output (traffic exiting the interface) or Input (traffic entering the interface). Queuing and shaping can be applied only to output traffic.

Step 6 (Optional) Define interface-level (hierarchical) shaping parameters. See Table 66-4 on page 66-21 for details.

Note When you enable hierarchical shaping on an interface, you cannot define shaping parameters for specific QoS classes. Shaping can be used only on output traffic. See Understanding Policing and Shaping Parameters, page 66-6 for more information about shaping.

Step 7 Click **OK**. The QoS interface definition is displayed in the upper table of the Quality of Service page.

Note To edit a QoS interface definition, select an interface from the upper table, then click the **Edit** button. To remove an interface definition, select it from the table, then click the **Delete** button. You cannot delete an interface that has defined classes.

Step 8 With the interface selected in the upper table, click the **Add** button beneath the QoS Classes table. The QoS Class dialog box is displayed. See Table 66-5 on page 66-23 for a description of the fields in this dialog box.

The QoS Class dialog box enables you to determine which traffic over the selected interface is included in the QoS class and how to handle that traffic.

Step 9 (Optional) Select the **Default class** check box if you are defining the properties of the default QoS class for this interface. The default class is assigned to all traffic that does not match the criteria of the other defined classes.

Step 10 Define the QoS class using one or more tabs in the QoS Class dialog box, as described in:

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS Class Shaping Parameters, page 66-18

Step 11 Repeat Step 8 through Step 10 to add QoS classes to the interface defined in Step 3. If required, use the **Up Row** and **Down Row** buttons to reorder the classes.

Note To edit a QoS class, select the relevant interface from the upper table to display its defined classes in the QoS Class table. Select the class to edit, then click the **Edit** button. To remove a class, select it from the table, then click the **Delete** button.

Step 12 Repeat Step 3 through Step 11 to define QoS classes for a different interface on the selected router.

Defining QoS on the Control Plane

When you configure QoS on input traffic entering the control plane, you can define multiple QoS classes, including a default class for traffic that does not match the criteria you define for the other classes. After defining the matching criteria for a particular class, you can configure a policing definition for that class. (Marking, queuing, and shaping are not available.) For more information, see Understanding Control Plane Policing, page 66-9.

QoS policies defined on the control plane override any QoS parameters defined on an interface of the same device.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Before You Begin

Ensure that Cisco Express Forwarding (CEF) is enabled on the router. For more information, see CEF Interface Settings on Cisco IOS Routers, page 62-25.

Related Topics

- Defining QoS Policies, page 66-10
- Defining QoS on Interfaces, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Select an existing policy or create a new one.

The Quality of Service page is displayed. See Table 66-3 on page 66-20 for a description of the fields on this page.

Step 2 In the Applied to field, select **Control Plane** to define QoS policing on input traffic entering the control plane.

Step 3 Click the **Add** button beneath the Control Plane QoS Classes table. The QoS Class dialog box is displayed. See Table 66-5 on page 66-23 for a description of the fields in this dialog box.

The QoS Class dialog box enables you to determine which traffic over the selected interface is included in the QoS class and how to handle that traffic.

Step 4 (Optional) Select the **Default class** check box if you are defining the properties of the default QoS class for the control plane. The default class is assigned to all traffic that does not match the criteria of the other defined classes.

Step 5 Define the QoS class using the tabs in the QoS Class dialog box, as described in the following sections:

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Policing Parameters, page 66-17

Step 6 Repeat Step 3 through Step 5 to add QoS classes to the control plane. If required, use the **Up Row** and **Down Row** buttons to reorder the classes.

Defining QoS Class Matching Parameters

When you define matching parameters, you must define matching criteria and specify whether packets must meet one or all of the criteria to be considered part of the class. See Understanding Matching Parameters, page 66-2 for more information.

Note You do not define matching parameters when configuring the default class.

Related Topics

- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS Class Shaping Parameters, page 66-18
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Matching** tab. See Table 66-6 on page 66-24 for a description of the fields on this tab.

Step 3 Select a matching method:

- Any—Traffic matching any of the defined parameters is included in this class.
- All—Only traffic matching all of the defined parameters is included in this class.

Step 4 (Optional) Under Protocol, click **Add** to display a selector for choosing the protocols to include in this class. Select one or more items from the Available Protocols list, then click **>>** to add them to the Selected Protocols list.

Note When configuring QoS on the control plane, only the ARP protocol can be selected.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the Protocol field.

Step 5 (Optional) Under Precedence, click **Add** to display a selector for choosing which IP precedence values (from 0 to 7) to include in this class. Select one or more items from the Available Precedences list, then click **>>** to add them to the Selected Precedences list. Traffic that arrives marked with one of these values matches this criterion.

Note For more information about IP precedence values, see Table 66-1 on page 66-3.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the Precedences field.

Step 6 (Optional) Under DSCP, click **Add** to display a selector for choosing which DSCP values (from 0 to 63) to include in this class. Select one or more items (up to eight) from the Available DSCPs list, then click **>>** to add them to the Selected DSCPs list. Traffic that arrives marked with one of these values matches this criterion.

When you finish, click **OK** to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the DSCP field.

Step 7 (Optional) Under ACL, define ACLs as part of the matching criteria for this class:

a. Click **Edit** to display the Edit ACLs dialog box. Use this dialog box to define which ACLs to include in this class.

b. Enter one or more ACLs, or click **Select** to select an ACL object from a list or to create a new one. Traffic that matches these ACL definitions matches this criterion.

c. When you finish, click **OK** twice to save your definitions and return to the QoS Class dialog box. Your selections are displayed in the ACL field.

Tip Use the up and down arrows to order the ACLs. We recommend placing more frequently used ACLs at the top of the list to optimize the matching process.

Step 8 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 9 Do one of the following:

- When defining QoS on interfaces, continue as described in Defining QoS on Interfaces, page 66-10.
- When defining control plane policing, continue as described in Defining QoS on the Control Plane, page 66-12.

Defining QoS Class Marking Parameters

When you define marking parameters, you can mark the packets in this QoS class with either a precedence value or a DSCP value. See Understanding Marking Parameters, page 66-3 for more information.

Note Marking is not available when you configure QoS on the control plane.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS Class Shaping Parameters, page 66-18
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Marking** tab. See Table 66-7 on page 66-26 for a description of the fields on this tab.

Step 3 Select the **Enable Marking** check box.

Step 4 Select one of the following marking options:

- Precedence—Select an IP precedence value (0 to 7) from the displayed list. For more information about these values, see Table 66-1 on page 66-3.
- DSCP—Select a DSCP value (0 to 63) from the displayed list.

Step 5 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 6 Continue as described in Defining QoS Policies, page 66-10.

Defining QoS Class Queuing Parameters

When you define queuing parameters, you can specify the amount of available bandwidth to provide to the traffic in this QoS class. You can also define a fixed amount of bandwidth that must be provided to high-priority traffic; you can define the priority parameter on only one class per interface. In addition, you must specify the type of queue management to perform on this class. See Understanding Queuing Parameters, page 66-4 for more information.

Note Queuing is not available when you configure QoS on the control plane.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS Class Shaping Parameters, page 66-18
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Queuing and Congestion Avoidance** tab. See Table 66-8 on page 66-27 for a description of the fields on this tab.

Step 3 Click the **Enable Queuing and Congestion Avoidance** check box.

Queuing options depend on whether you are defining the default class or a different class:

- When you define any class other than the default class, select one of the following queuing options:
  - Priority—Define the amount of bandwidth to make available to high-priority traffic. Low-Latency Queuing, page 66-5 (LLQ) ensures that this traffic receives this fixed amount of bandwidth at all times. This is particularly useful for voice traffic, which requires low latency. You can define this amount by percentage or by an absolute value of kilobits per second.

    Note You can define this option for only one class per interface.

  - Bandwidth—Enter the amount of bandwidth to allocate to this class. You can define this amount by percentage or by an absolute value of kilobits per second.

    Note The sum of all class bandwidth allocations on an interface cannot exceed 100 percent of the total available bandwidth.

- When you define the default class, select one of the following queuing options:

- Fair queue—Enter the number of queues to reserve for the default class. Values range in powers of 2 from 16 to 4096. By default, the number of queues is based on the available bandwidth of the selected interface. For more information, seeTable 66-2 on page 66-6.
- Bandwidth—Enter the amount of bandwidth to allocate to this class. You can define this amount by percentage or by an absolute value of kilobits per second.

Step 4 (Optional) Define *one* of the following queue length management options:

- Queue Limit—(Default) Specify the maximum number of packets allowed. If you select this option, tail drop drops excess packets when the queue reaches its capacity.
- WRED Weight for Mean Queue Depth—WRED proactively drops packets until the transmitting protocol (usually TCP) responds by dropping its transmission rate, thereby alleviating congestion. Configure WRED by entering an exponential weight factor that is used to calculate the average queue size.

For more information, see Tail Drop vs. WRED, page 66-4.

Note You should change the default only if you are certain that your applications will benefit from a different value.

Note Do not use WRED with protocols that are not sufficiently robust to reduce their transmission rates in response to packet loss, such as IPX or AppleTalk. WRED cannot be configured when you select the Priority percent option.

Step 5 Go to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 6 Continue as described in Defining QoS Policies, page 66-10.

Defining QoS Class Policing Parameters

When you define policing parameters, you must specify the average data rate, which determines the amount of traffic that can be transmitted. In addition, you must specify the action to take on traffic bursts that exceed this data rate.

You can configure policing for all QoS classes, including the default class. For more information about policing, see Understanding Policing and Shaping Parameters, page 66-6.

You can also configure policing on the control plane. For more information, see Understanding Control Plane Policing, page 66-9.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS Class Shaping Parameters, page 66-18
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Policing** tab. See Table 66-9 on page 66-29 for a description of the fields on this tab.

Step 3 Select the **Enable Policing** check box.

Step 4 Define CIR, confirm burst, and excess burst values. You can define the CIR by percentage or by an absolute value of bits per second. The option you choose determines how you define the burst values.

Step 5 Select the action to perform on packets that conform to the rate limit:

- transmit—Transmit the packet.
- set-prec-transmit—Set the IP precedence to a defined value, then send the packet. This option is not available when configuring QoS on the control plane.
- set-dscp-transmit—Set the DSCP to a defined value, then send the packet. This option is not available when configuring QoS on the control plane.
- drop—Drop the packet.

Step 6 Select the action to perform on exceed packets. The list of available actions depends on the selected conform action.

For example, if transmit is performed on conforming packets, you can select any of the actions listed in Step 5 for exceeding packets. However, if you selected one of the set actions for conforming packets, you can select only a set action or the drop action for exceeding packets. If you selected drop as the conform action, you must select drop as the exceed action.

Step 7 Select the action to perform on violate packets. The list of available actions depends on the selected exceed action.

For example, if transmit is performed on exceeding packets, you can select any of the actions listed in Step 5 for violating packets. However, if you selected one of the set actions for exceeding packets, you can select only a set action or the drop action for violating packets. If you selected drop as the exceed action, you must select drop as the violate action.

Step 8 Go to another tab, or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 9 Do one of the following:

- When defining QoS on interfaces, continue as described in Defining QoS Policies, page 66-10.
- When defining control plane policing, continue as described in Defining QoS on the Control Plane, page 66-12.

Defining QoS Class Shaping Parameters

When you define shaping parameters, you must specify whether to base traffic shaping on the average data rate or on the average data rate plus the excess burst rate that occurs during traffic peaks. In both cases, traffic that exceeds these definitions is buffered until the rate lowers, allowing the packets to be sent.

The following conditions pertain:

- Shaping can be used only on output traffic.
- Shaping can be configured for all QoS classes, including the default class.
- Shaping is not available when you configure the QoS class for priority traffic.

- Shaping is not available when you configure QoS on the control plane.

For more information about shaping, see Understanding Policing and Shaping Parameters, page 66-6.

Tip To configure shaping on all the QoS classes defined for the interface (hierarchical shaping), see Defining QoS on Interfaces, page 66-10.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1

Step 1 On the Quality of Service page, click the **Add** button beneath the QoS Classes table, or select a class and then click the **Edit** button. The QoS Class dialog box is displayed.

Step 2 Click the **Shaping** tab. See Table 66-10 on page 66-31 for a description of the fields on this tab.

Step 3 Select the **Enable Shaping** check box.

Step 4 Select the shaping type (Average or Peak).

Step 5 Define CIR, sustained burst, and excess burst values. You can define the CIR by percentage or by an absolute value of bits per second. The option you choose determines how you define the burst values.

Step 6 Proceed to another tab or click **OK** to save your definitions locally on the client and close the dialog box. The defined class is displayed in the QoS Classes table on the Quality of Service page.

Step 7 Continue as described in Defining QoS Policies, page 66-10.

Quality of Service Policy Page

Use the Quality of Service page to view, create, and edit QoS classes on specific interfaces of the selected device or on the control plane. QoS policies enable you to define techniques for managing the delay, delay variation (jitter), bandwidth, and packet loss parameters on a network. In addition, you can use the Quality of Service page to configure hierarchical shaping on an interface as an alternative to configuring shaping parameters for individual QoS classes.

For more information, see Quality of Service on Cisco IOS Routers, page 66-1.

Navigation Path

- (Device view) Select **Platform > Quality of Service** from the Policy selector.
- (Policy view) Select **Router Platform > Quality of Service** from the Policy Type selector. Create a new policy or select an existing policy from the Shared Policy selector.

Related Topics

- Defining QoS Policies, page 66-10

- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 66-3 Quality of Service Page

Element	Description
Apply To	The router component on which to define the QoS policy: • Interfaces—Configures QoS classes on specific interfaces. • Control Plane—Configures QoS on the router control plane. See Understanding Control Plane Policing, page 66-9. **Note** If you configure QoS on both the interfaces and the control plane of the same device, only the control plane configuration is deployed.
Interface Table	If you are defining classes on interfaces, the upper table lists the interfaces on which you are defining QoS classes. The direction column indicates the direction of traffic through the interface to which the classes apply (Output or Input). The classes you can define vary based on the direction. The other fields indicate whether you defined shaping on the interface, and if shaping is defined, the type of hierarchical shaping (average or peak), the committed information rate (CIR), and the sustained and excess burst size. For detailed information about the attributes, see QoS Policy Dialog Box, page 66-21. • To add an interface to the table, click the Add button. • To edit the settings for an interface, select it and click the Edit button. • To delete an interface, select it and click the Delete button.
QoS Classes Table	The classes defined for the interface selected in the upper table, or for the control plane. Each row represents a separate class. The No. column indicates the order of the classes, and is very important: QoS is applied to packets on a first-match basis, based on class order. The Default Class column indicates whether this class is the default for all packets on the interface that do not match the criteria of the other defined classes. Make this the last class in the list. The remaining columns indicate the match criteria for the class, and the packet marking, queuing and congestion avoidance, policing, and shaping defined for the class, if any. For detailed information about the attributes, see QoS Class Dialog Box, page 66-23. • To add class to the table, click the Add button. • To edit the settings for a class, select it and click the Edit button. • To delete a class, select it and click the Delete button. • To change the order of a class, select it and click the Up and Down arrow buttons to reposition it.

QoS Policy Dialog Box

Use the QoS Policy dialog box to select an interface on which you want to define QoS parameters. In addition, you can use this dialog box to configure a single set of shaping parameters for all the traffic on the selected interface (known as hierarchical shaping). Using hierarchical shaping eliminates the need to configure shaping parameters for each QoS class defined on the interface.

Note This dialog box is not applicable when defining a QoS policy on the control plane. For more information, see Defining QoS on the Control Plane, page 66-12.

After you create your QoS interface definitions, you can define one or more QoS classes for each interface. For more information, see QoS Class Dialog Box, page 66-23.

Navigation Path

Go to the Quality of Service Policy Page, page 66-19, then click the **Add** or **Edit** button beneath the upper table to define a QoS interface definition.

Related Topics

- Defining QoS Policies, page 66-10
- Quality of Service on Cisco IOS Routers, page 66-1
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 66-4 QoS Policy Dialog Box

Element	Description
Interface	The interface on which QoS is defined. Enter the name of an interface or interface role, or click **Select** to select an object from a list or to create a new object.
Direction	The direction of the traffic on which to configure QoS: • Output—Traffic that exits the interface. • Input—Traffic that enters the interface.
Hierarchical Shaping settings	
Enable Shaping	When selected, configures hierarchical traffic shaping on the selected interface. When deselected, hierarchical shaping is not used. **Note** Shaping can be performed only on output traffic.

Table 66-4 QoS Policy Dialog Box (continued)

Element	Description
Type	The type of shaping to perform: • Average—Limits the data rate for each interval to the sustained burst rate (also known as the Committed Burst rate or Bc), achieving an average rate no higher than the committed information rate (CIR). Additional packets are buffered until they can be sent. • Peak—Limits the data rate for each interval to the sustained burst rate plus the excess burst rate (Be). Additional packets are buffered until they can be sent.
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 1000000000 bits per second, and must be in multiples of 8000. Although data bursts during an interval may exceed this rate, the average data rate over any multiple integral of the interval will not exceed this rate.
Sustained Burst	The normal burst size. If you select average as the shaping type, data bursts during an interval are limited to this value. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** We recommend that you leave this field blank when the CIR is defined by an absolute value. This allows the algorithms used by the device to determine the optimal sustained burst value.
Excess Burst	The excess burst size. If you select peak as the shaping type, data bursts during an interval can equal the sum of the sustained burst value plus this value. The average data rate over multiple intervals, however, will continue to conform to the CIR. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** If you do not configure this field when the CIR is defined by an absolute value, the sustained burst value is used.

QoS Class Dialog Box

Use the QoS Class dialog box to create or edit a QoS class on a selected interface or control plane of a Cisco IOS router. You can define up to 16 classes on a single interface and 256 classes for the device as a whole.

Note QoS is applied to packets on a first-match basis. The router examines the table of QoS classes starting from the top and applies the properties of the first class whose matching criteria matches the packet. Therefore, it is important that you define and order your classes carefully. The default class should be placed last to prevent traffic that matches a specific class from being treated as unmatched traffic.

Navigation Path

Go to the Quality of Service Policy Page, page 66-19. Complete the options at the top of the page, then do one of the following:

- To create a QoS class, select an interface from the upper table, then click the **Add** button beneath the QoS Class table. When creating a QoS class for the control plane, just click the **Add** button beneath the table.
- To edit a QoS class:
  - Select the interface whose class you want to edit from the upper table (Not required when selecting the control plane.).
  - Select the relevant class defined for that interface in the QoS Classes table. (Not required when selecting the control plane.)
  - Click the **Edit** button under the QoS Class table.

Related Topics

- QoS Policy Dialog Box, page 66-21
- Defining QoS Policies, page 66-10
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12

Field Reference

Table 66-5 QoS Class Dialog Box

Element	Description
Set as Default Class	When selected, enables you to define the default class for all traffic that does not match the other QoS classes on this interface. When deselected, enables you to define a specific QoS class on this interface. **Note** When you define the default class, you do not configure any matching parameters; by definition the class consists of all traffic that does not match any of the other classes. Therefore, the Matching tab is disabled.
Matching tab	Defines the traffic that is included in this QoS class. See QoS Class Dialog Box—Matching Tab, page 66-24.

Table 66-5 QoS Class Dialog Box (continued)

Element	Description
Marking tab	Marks the traffic in this class so that downstream devices can properly identify it. See QoS Class Dialog Box—Marking Tab, page 66-26.
Queuing and Congestion Avoidance tab	Defines how to queue the output traffic in this class. See QoS Class Dialog Box—Queuing and Congestion Avoidance Tab, page 66-27.
Policing tab	Limits the traffic flow for this class to a configured rate. See QoS Class Dialog Box—Policing Tab, page 66-29.
Shaping tab	Controls the flow of output traffic for this class so that it conforms with the requirements of downstream devices. See QoS Class Dialog Box—Shaping Tab, page 66-31.

Note When you configure a QoS policy on the control plane, only the Matching tab and Policing tab are available.

QoS Class Dialog Box—Matching Tab

Use the Matching tab of the QoS Class dialog box to define which traffic over the selected interface is considered to be part of this class.

Note When you define the default class, the Matching tab is disabled.

Navigation Path

Go to the QoS Class Dialog Box, page 66-23, then click the **Matching** tab.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19
- Creating Access Control List Objects, page 6-53

Field Reference

Table 66-6 QoS Class Dialog Box—Matching Tab

Element	Description
Match Method	The traffic matching option used for this class: • Any—Assigns traffic matching any of the defined class map criteria to this QoS class. • All—Assigns only traffic matching all of the defined class map criteria to this QoS class.

Table 66-6 QoS Class Dialog Box—Matching Tab (continued)

Element	Description
Protocol	One or more protocols included in this class map. Click **Add** to display a selector. Select one or more items from the Available Protocols list, then click >> to add them to the Selected Protocols list. The only protocol available for the control plane is ARP; ARP and CDP are not available for input classes configured on an interface. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the Protocol field. **Note** To remove a protocol from the QoS class, select it from the Protocol field, then click **Delete**.
Precedence	One or more IP Precedence (IPP) values included in this class map. Click **Add** to display a selector. Select one or more items from the Available Precedences list, then click >> to add them to the Selected Precedences list. For more information about IP precedence values, see Table 66-1 on page 66-3. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the Precedence field. **Note** To remove an IPP value from the QoS class, select it from the Precedence field, then click **Delete**.
DSCP	One or more Differentiated Services Code Point (DSCP) values included in this class map. Click **Add** to display a selector. Select one or more items (up to eight) from the Available DSCPs list, then click >> to add them to the Selected DSCPs list. When you finish, click **OK** to return to the QoS Class dialog box. Your selections are displayed in the DSCP field. **Note** To remove a DSCP value from the QoS class, select it from the DSCP field, then click **Delete**.
ACL	The ACLs that are used for defining which traffic requires QoS. Click **Edit** to add or remove ACL objects. Use the up and down arrows to order the ACLs in the list. We recommend that you place frequently used ACLs at the top of the list to optimize the matching process.

Edit ACLs Dialog Box—QoS Classes

When configuring a QoS policy on a Cisco IOS router, use the Edit ACLs dialog box to specify which ACLs should be included in the matching criteria for the selected QoS class. Traffic matching this criteria is included as part of the class.

Enter the names of the extended ACLs or click **Select** to select an ACL object from a list or to create a new one. Separate multiple ACL objects with commas and place them in priority order.

For more information, see Creating Extended Access Control List Objects, page 6-54.

Navigation Path

Go to the QoS Class Dialog Box—Matching Tab, page 66-24, then click **Edit** in the ACL field.

Related Topics

- Defining QoS Class Matching Parameters, page 66-13
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19
- Selecting Objects for Policies, page 6-2

QoS Class Dialog Box—Marking Tab

Use the Marking tab of the QoS Class dialog box to classify packets. Traffic policers and shapers use these classifications to ensure adherence to the contracted level of service. Downstream devices use this classification to identify the packets and apply the appropriate QoS functions to them.

Note The Marking tab is unavailable when you define a QoS policy on the control plane.

Navigation Path

Go to the QoS Class Dialog Box, page 66-23, then click the **Marking** tab.

Related Topics

- Defining QoS Class Marking Parameters, page 66-15
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19

Field Reference

Table 66-7 QoS Class Dialog Box—Marking Tab

Element	Description
Enable Marking	When selected, enables you to mark the traffic in this QoS class with a specific precedence or DSCP value (regardless of any value the traffic might have had when it first entered the device). This mark enables downstream devices to identify the traffic and apply the appropriate QoS features to it. When deselected, disables all marking options for the selected QoS class. The traffic in this QoS class maintains its original precedence or DSCP value, if any.

Table 66-7 QoS Class Dialog Box—Marking Tab (continued)

Element	Description
Precedence	The precedence value with which to mark the traffic in this class: • network (7) • internet match (6) • critical (5) • flash-override (4) • flash (3) • immediate (2) • priority (1) • routine (0)
DSCP	The DSCP value (0 to 63) with which to mark the traffic in this class.

QoS Class Dialog Box—Queuing and Congestion Avoidance Tab

Use the Queuing and Congestion Avoidance tab of the QoS Class dialog box to perform Class-Based Weighted Fair Queuing (CBWFQ) on the output traffic in the selected QoS class. Queuing prioritizes traffic and manages congestion on your network by determining the order in which packets are sent out over an interface. Queuing and congestion avoidance applies only to interface classes for output traffic.

The fields displayed in the Queuing tab depend on whether you are defining a specific QoS class or the default class (by selecting **Set as Default Class**), and also by the type of router and the Cisco IOS software version.

Navigation Path

Go to the QoS Class Dialog Box, then click the **Queuing and Congestion Avoidance** tab.

Related Topics

- Defining QoS Class Queuing Parameters, page 66-16
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19

Field Reference

Table 66-8 QoS Class Dialog Box—Queuing and Congestion Avoidance Tab

Element	Description
Enable Queuing and Congestion Avoidance	Whether to configure queuing and congestion avoidance properties in the QoS class.

Table 66-8 QoS Class Dialog Box—Queuing and Congestion Avoidance Tab (continued)

Element	Description
Priority (Non-default classes only.)	Configure low-latency queuing (LLQ) in this class to ensure that priority traffic, such as voice traffic, receives the defined bandwidth (see Low-Latency Queuing, page 66-5). Specify the amount of bandwidth allocated to high-priority traffic on this interface by: • Percentage—Valid values range from 1 to 100%. • Kbit/sec—Valid values range from 8-2000000 kilobits per second. **Note** You can define this option for one class only per interface. If you select this option, the Shaping tab is disabled.
Fair Queue Number of Dynamic Queues (Default class only.)	Configure class-based weighted fair queuing in this class. If the device is running an IOS software version lower than 12.4(20)T, you must specify the number of dynamic queues to reserve for this class. You should base your number on the available bandwidth of the interface. You can specify a number between 16 and 4096 that is a power to 2. For information on the default number of queues the device uses, see Default Class Queuing, page 66-6. Available bandwidth is evenly distributed among the queues unless you configure a queue limit. **Tip** Failure to provide a sufficient number of queues for the default class (a condition known as starvation) could result in the traffic not being sent.
Bandwidth	Configure the minimum bandwidth to guarantee to this class. You can define this amount by: • Percentage—Valid values range from 1 to 100% of the total available bandwidth. • Kbit/sec—Valid values range from 8-2000000 kilobits per second.
Enable Fair Queue (Non-default class only.)	When you configure bandwidth for a non-default class, whether to also enable class-based weighted fair queuing (CBWFQ). The device calculates the number of queues to configure based on the available bandwidth, and distributes the bandwidth evenly among the queues unless you configure a queue limit. This option is available only for Aggregation Services Routers (ASR) and for routers running 12.4(20)T and higher.
Queue Limit	The maximum number of packets that can be queued for the class. Any additional packets are dropped using tail drop until the congestion is gone. This is the default option for limiting queue size unless Weighted Random Early Detection (WRED) is configured.

Table 66-8 QoS Class Dialog Box—Queuing and Congestion Avoidance Tab (continued)

Element	Description
WRED Weight for Mean Queue Depth	The exponential weight factor to use to calculate the average queue size. Use this option when defining WRED instead of tail drop (queue limit) for this class. When the queue size exceeds the value determined by this weight factor, WRED randomly discards packets until the transmitting protocol decreases its transmission rate to ease congestion. Exponent values range from 1 to 16. The default is 9. This option is best suited for protocols like TCP, which respond to dropped packets by decreasing the transmission rate. We recommend that you do not change the default unless you determine that your applications would benefit from the change.

QoS Class Dialog Box—Policing Tab

Use the Policing tab of the QoS Class dialog box to configure rate limits on the traffic in a selected QoS class. Excess traffic is either dropped or transmitted with a different (typically lower) priority.

Navigation Path

Go to the QoS Class Dialog Box, page 66-23, then click the **Policing** tab.

Related Topics

- Defining QoS Class Policing Parameters, page 66-17
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19

Field Reference

Table 66-9 QoS Class Dialog Box—Policing Tab

Element	Description
Enable Policing	When selected, enables you to configure Class-Based Policing to control the maximum rate of traffic for this class. Security Manager uses a two-token bucket algorithm, which includes a defined violate action that is performed when neither bucket can accommodate the incoming packet. When deselected, disables all policing options for the selected QoS class.

Table 66-9 QoS Class Dialog Box—Policing Tab (continued)

Element	Description
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 2000000000 bits per second. In the token bucket algorithm, this rate represents the token arrival rate for filling both token buckets. Traffic that falls under this rate always conforms. **Note** When you configure Understanding Control Plane Policing, page 66-9, you must define the CIR in bits per second.
Conform Burst	The normal burst size, which determines how large traffic bursts can be before some traffic exceeds the rate limit. In the token bucket algorithm, it represents the full size of the first (conform) token bucket. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 1 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000-512000000 bytes.
Excess Burst	The excess burst size, which determines how large traffic bursts can be before all traffic exceeds the rate limit. In the token bucket algorithm, it represents the full size of the second (exceed) token bucket. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 1 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000-512000000 bytes.
Conform action	The action to take on packets that conform to the rate limit: • transmit—Transmits the packet. • set-prec-transmit—Sets the IP precedence to a value you specify (0 to 7) and then sends the packet. Not available on the control plane. • set-dscp-transmit—Sets the DSCP to a value you specify (0 to 63) and then sends the packet. Not available on the control plane. • drop—Drops the packet.
Exceed action	The action to take on packets that exceed the rate limit, but can be handled using the second (exceed) token bucket. The actions available for selection depend on the defined conform action. For example, if you select one of the set options as the conform action, you cannot select transmit as the exceed action. If you select drop as the conform action, then you must also select drop as the exceed action.

Table 66-9 QoS Class Dialog Box—Policing Tab (continued)

Element	Description
Violate action	The action to take on packets that cannot be serviced by either the conform bucket or the exceed bucket. The actions available for selection depend on the defined exceed action. For example, if you select one of the set options as the exceed action, you cannot select transmit as the violate action. If you select drop as the exceed action, then you must also select drop as the violate action.

QoS Class Dialog Box—Shaping Tab

Use the Shaping tab of the QoS Class dialog box to control the rate of output traffic for the selected QoS class. Shaping typically delays excess traffic by using a buffer, or queuing mechanism, to hold packets and shape the flow when the data rate of the source is higher than expected.

Note The Shaping tab is unavailable when you define a QoS policy on the control plane, use hierarchical shaping on the interface, define a QoS class for input traffic, or perform queuing on priority traffic.

Navigation Path

Go to the QoS Class Dialog Box, page 66-23, then click the **Shaping** tab.

Related Topics

- Defining QoS Class Shaping Parameters, page 66-18
- Defining QoS on Interfaces, page 66-10
- Defining QoS on the Control Plane, page 66-12
- Quality of Service Policy Page, page 66-19

Field Reference

Table 66-10 QoS Class Dialog Box—Shaping Tab

Element	Description
Enable Shaping	When selected, enables you to configure Distributed Traffic Shaping (DTS) to control the rate of traffic for this class. DTS uses queues to buffer traffic surges that can congest the network. When deselected, disables all shaping options for the selected QoS class. **Note** Shaping can be performed only on output traffic.

Table 66-10 QoS Class Dialog Box—Shaping Tab (continued)

Element	Description
Type	The type of shaping to perform: • Average—Limits the data rate for each interval to the sustained burst rate (also known as the committed burst rate or Bc), achieving an average rate no higher than the committed information rate (CIR). Additional packets are buffered until they can be sent. • Peak—Limits the data rate for each interval to the sustained burst rate plus the excess burst rate (Be). Additional packets are buffered until they can be sent.
CIR	The average data rate (also known as the committed information rate or CIR). You can define this amount by: • Percentage—Valid values range from 0 to 100% of the overall available bandwidth. • Bit/sec—Valid values range from 8000 to 1000000000 bits per second, and must be in multiples of 8000. Although data bursts during an interval may exceed this rate, the average data rate over any multiple integral of the interval will not exceed this rate.
Sustained Burst	The normal burst size. If you select average as the shaping type, data bursts during an interval are limited to this value. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** We recommend that you leave this field blank when the CIR is defined by an absolute value. This allows the algorithms used by the device to determine the optimal sustained burst value.
Excess Burst	The excess burst size. If you select peak as the shaping type, data bursts during an interval can equal the sum of the sustained burst value plus this value. The average data rate over multiple intervals, however, will continue to conform to the CIR. The range of valid values is determined by the CIR: • When the CIR is defined by percentage—Valid values range from 10 to 2000 milliseconds. • When the CIR is defined by an absolute value—Valid values range from 1000 to 154400000 bytes, in multiples of 128 bytes. **Note** If you do not configure this field when the CIR is defined by an absolute value, the sustained burst value is used.

Configuring Routing Policies

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

This chapter contains the following topics:

- BGP Routing on Cisco IOS Routers, page 67-1
- BGP Routing Policy Page, page 67-4
- EIGRP Routing on Cisco IOS Routers, page 67-8
- EIGRP Routing Policy Page, page 67-13
- OSPF Routing on Cisco IOS Routers, page 67-19
- OSPF Interface Policy Page, page 67-30
- OSPF Process Policy Page, page 67-34
- RIP Routing on Cisco IOS Routers, page 67-42
- RIP Routing Policy Page, page 67-45
- Static Routing on Cisco IOS Routers, page 67-50
- Static Routing Policy Page, page 67-51

BGP Routing on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

BGP is an Exterior Gateway Protocol (EGP) that guarantees the loop-free exchange of routing information between autonomous systems (ASs). The *primary* function of a BGP system is to exchange information with other BGP systems about the networks it can reach, including AS path information. This information can be used to construct a graph of AS connectivity from which routing loops can be pruned and with which AS-level policy decisions can be enforced.

BGP is the routing protocol used on the Internet and is commonly used between Internet service providers. To achieve scalability at this level, BGP uses several route parameters (attributes) to define routing policies and maintain a stable routing environment. Additionally, BGP uses classless interdomain routing (CIDR) to greatly reduce the size of Internet routing tables.

A BGP route consists of a network number, a list of ASs through which information has passed (called the *autonomous system path*), and the defined path attributes.

A BGP router exchanges routing information only with those routers that you define as its neighbors. BGP neighbors exchange complete routing information when the TCP connection between them is established. Updates are sent to neighbors only when changes to the routing table are detected. BGP routers do not send regular, periodic updates.

The following topics describe the tasks you perform to create a BGP routing policy:

- Defining BGP Routes, page 67-2
- Redistributing Routes into BGP, page 67-3

Note Security Manager supports versions 2, 3 and 4 of BGP, as defined in RFCs 1163, 1267 and 1771.

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- RIP Routing on Cisco IOS Routers, page 67-42
- OSPF Routing on Cisco IOS Routers, page 67-19
- EIGRP Routing on Cisco IOS Routers, page 67-8

Defining BGP Routes

As with all EGPs, when you configure a BGP routing policy, you must define the relationship the router has with its neighbors. BGP supports two kinds of neighbors: internal (located in the same AS) and external (located in a different AS). Typically, external neighbors are adjacent to each other and share a subnet; internal neighbors can be anywhere in the same AS.

In addition, you can select whether to enable the following optional features:

- Auto-summarization
- Synchronization
- Neighbor logging

If enabled, *auto-summarization* injects only the network route when a subnet is redistributed from an Interior Gateway Protocol (IGP) such as OSPF or EIGRP into BGP. *Synchronization* is useful if your AS acts as an intermediary, passing traffic from one AS to another AS, because it ensures that your AS is consistent about the routes it advertises. For example, if BGP were to advertise a route before all routers in your network had learned about the route through your IGP, your AS might receive traffic that some routers cannot yet route. *Neighbor logging* enables the router to keep track of messages issued by BGP neighbors when they reset, become unreachable, or restore their connection to the network.

This procedure describes how to define a BGP route. You can define only one BGP route on each router.

Related Topics

- Redistributing Routes into BGP, page 67-3
- BGP Routing on Cisco IOS Routers, page 67-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > BGP** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Routing > BGP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The BGP Setup is displayed. See Table 67-1 on page 67-5 for a description of the fields on this tab.

Step 2 On the BGP Setup tab, enter the AS number to which the router belongs.

Step 3 (Optional) Enter the addresses of the networks that are local to this AS. You can use a combination of addresses and network/host objects, or click **Select** to select an object from a list or to create a new one. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 4 Define external and internal BGP neighbors for the routers:

a. Click **Add** under Neighbors to display the BGP Neighbors dialog box. See Table 67-2 on page 67-6 for a description of the fields in this dialog box.

b. Enter an AS number and then click **Select** to select the hosts that are neighbors within the defined AS. Internal neighbors are located in the same AS as the router; external neighbors are located in a different AS.

c. Click **OK** to save your definitions and return to the BGP Neighbors dialog box.

d. (Optional) Repeat b. and c. to define neighbors in additional ASs.

Note When you define BGP neighbors, the IP addresses cannot belong to an interface on the selected router. In addition, you cannot define the same IP address in more than one AS.

When you finish, click **OK** in the BGP Neighbors dialog box to return to the BGP Setup tab. Your selections are displayed in the Neighbors field.

Step 5 (Optional) Select the Auto-Summary check box to enable automatic summarization. If automatic summarization is enabled, only the network route is injected into the BGP table when a subnet is redistributed from an IGP (such as OSPF or EIGRP) into BGP.

Step 6 (Optional) Select the **Synchronization** check box to synchronize BGP with the IGP. Enabling this feature causes BGP to wait until the IGP propagates routing information across the AS.

You do not need synchronization if your AS does not pass traffic it receives from one AS to another AS, or if all the routers in your AS run BGP. Disabling synchronization enables BGP to converge more quickly.

Step 7 (Optional) Select the **Log-Neighbor** check box to enable the logging of messages generated when a BGP neighbors resets, comes up, or goes down.

Redistributing Routes into BGP

Redistribution refers to using a routing protocol, such as BGP, to advertise routes that are learned by some other means, such as a different routing protocol, static routes, or directly connected routes. For example, you can redistribute routes from the OSPF routing protocol into your BGP autonomous system (AS). Redistribution is necessary in networks that operate in multiple-protocol environments and can be applied to all IP-based routing protocols.

Before You Begin

- Define a BGP AS. See Defining BGP Routes, page 67-2.

Related Topics

- Defining BGP Routes, page 67-2
- BGP Routing on Cisco IOS Routers, page 67-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > BGP** from the Policy selector, then click the **Redistribution** tab in the work area.
- (Policy view) Select **Router Platform > Routing > BGP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Redistribution** tab.

The BGP Redistribution tab is displayed. See Table 67-3 on page 67-7 for a description of the fields on this tab.

Step 2 On the BGP Redistribution tab, select a row from the BGP Redistribution Mappings table, then click **Edit**, or click **Add** to create a mapping. The BGP Redistribution Mapping dialog box appears. See Table 67-4 on page 67-8 for a description of the fields in this dialog box.

Step 3 Select the protocol whose routes you want to redistribute into BGP.

Note You can create a single mapping for each static route, RIP route, EIGRP AS, and OSPF process.

Step 4 (Optional) Modify the default metric (cost) of the redistributed routes. The metric determines the priority of the routes.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box. The redistribution mapping appears in the Redistribution Mapping table in the BGP Redistribution tab.

BGP Routing Policy Page

Border Gateway Protocol (BGP) is an exterior gateway protocol (EGP) that performs routing between multiple autonomous systems or domains and exchanges routing and reachability information with other BGP systems. BGP is used to exchange routing information on the Internet and is the protocol used between Internet service providers.

You can configure BGP routing policies from the following tabs on the BGP Routing page:

- BGP Page—Setup Tab, page 67-5
- BGP Page—Redistribution Tab, page 67-7

For more information, see BGP Routing on Cisco IOS Routers, page 67-1.

Navigation Path

- (Device view) Select **Platform > Routing > BGP** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > BGP** from the Policy Type selector. Right-click **BGP** to create a policy, or select an existing policy from the Shared Policy selector.

BGP Page—Setup Tab

Use the BGP Setup tab to define the number of the autonomous system (AS) in which the selected router is located. You must then define which networks are included in the AS and which networks are the internal and external neighbors of the router. Additionally, you can enable or disable options that govern the interaction between BGP and Interior Gateway Protocols (IGPs), such as OSPF and EIGRP. Use a third option to enable the logging of messages from BGP neighbors.

Navigation Path

Go to the BGP Routing Policy Page, page 67-4, then click the **Setup** tab.

Related Topics

- Defining BGP Routes, page 67-2
- BGP Page—Redistribution Tab, page 67-7
- Specifying IP Addresses During Policy Definition, page 6-87
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 67-1 BGP Setup Tab

Element	Description
AS Number	The number of the autonomous system in which the router is located. Valid values range from 1 to 65535. This number enables a BGP routing process. If BGP is already configured on the device, you cannot successfully change and deploy this number. If you need to change the AS number, first unassign the BGP policy, deploy your change (thus removing the BGP configuration from the device), then configure the BGP policy with the new number and redeploy the configuration.
Networks	The networks associated with the BGP route. Enter one or more network addresses or network/host objects, or click **Select** to select the object from a list or to create a new one. **Note** To remove a network from the route, select it from the Network field, then click **Delete**.
Neighbors	The *internal* neighbors (those located in the same AS as the router) and *external* neighbors (located in different ASs) of the router. See Neighbors Dialog Box, page 67-6.
Auto-Summary	When selected, automatic summarization is enabled. When a subnet is redistributed from an IGP (such as RIP, OSPF or EIGRP) into BGP, this BGP version 3 feature injects only the network route into the BGP table. Automatic summarization reduces the size and complexity of the routing table that the router must maintain. When deselected, automatic summarization is disabled. This is the default.

Table 67-1 BGP Setup Tab (continued)

Element	Description
Synchronization	When selected, synchronization is enabled. Use this feature to ensure that all routers in your network are consistent about the routes they advertise. Synchronization forces BGP to wait until the IGP propagates routing information across the AS. When deselected, synchronization is disabled. You can disable synchronization if this router does not pass traffic from a different AS to a third AS, or if all the routers in the AS are running BGP. Disabling this feature has the benefit of reducing the number of routes the IGP must carry, which improves convergence times. This is the default.
Log-Neighbor	When selected, enables the logging of messages that are generated when a BGP neighbors resets, connects to the network, or is disconnected. This is the default. When deselected, message logging is disabled.

Neighbors Dialog Box

Use the Neighbors dialog box to define the internal and external neighbors of the selected router.

Navigation Path

Go to the BGP Page—Setup Tab, page 67-5, then click the **Add** or **Edit** button in the Neighbors field.

Related Topics

- Defining BGP Routes, page 67-2
- Specifying IP Addresses During Policy Definition, page 6-87
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 67-2 Neighbors Dialog Box

Element	Description
AS Number	The number of the AS containing BGP neighbors. Internal neighbors have the same AS number as the network of the selected router. External neighbors have a different AS number.
IP Address	The IP addresses of the hosts that are neighbors of the router. BGP neighbors exchange routing information with each other whenever changes to the routing table are detected. When you define BGP neighbors, the IP addresses cannot belong to an interface on the selected router. In addition, you cannot define the same IP address in more than one AS. Enter one or more addresses or network/host objects, or click **Select** to select an object from a list or to create a new one. **Note** To remove a host from the list of BGP neighbors, select it from the Hosts field, then click **Delete**.

BGP Page—Redistribution Tab

Use the BGP Redistribution tab to view, create, edit, and delete redistribution settings when performing redistribution into a BGP autonomous system (AS).

Note You must define BGP setup parameters before you can access the BGP Redistribution tab. See BGP Page—Setup Tab, page 67-5.

Navigation Path

Go to the BGP Routing Policy Page, page 67-4, then click the **Redistribution** tab.

Related Topics

- Redistributing Routes into BGP, page 67-3
- BGP Page—Setup Tab, page 67-5
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-3 BGP Redistribution Tab

Element	Description
Protocol	The protocol that is being redistributed.
AS/Process ID	The AS number or process ID of the route being redistributed.
Metric	The value that determines the priority of the redistributed route.
Match	When redistributing an OSPF process, indicates the types of OSPF routes that are being redistributed.
Static Type	When redistributing static routes, indicates the type of static route, IP or OSI.
Add button	Opens the BGP Redistribution Mapping Dialog Box, page 67-7. From here you can define BGP redistribution mappings.
Edit button	Opens the BGP Redistribution Mapping Dialog Box, page 67-7. From here you can edit the selected BGP redistribution mapping.
Delete button	Deletes the selected BGP redistribution mappings from the table.

BGP Redistribution Mapping Dialog Box

Use the BGP Redistribution Mapping dialog box to add or edit the properties of a BGP redistribution mapping.

Navigation Path

Go to the BGP Page—Redistribution Tab, page 67-7, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Redistributing Routes into BGP, page 67-3

Field Reference

Table 67-4 BGP Redistribution Mapping Dialog Box

Element	Description
Protocol to Redistribute	The routing protocol that is being redistributed: • Static—Redistributes IP or OSI static routes. You can define a single mapping for each route. • EIGRP—Redistributes an EIGRP autonomous system. Enter the AS number in the displayed field. You can define a single mapping for each AS. • RIP—Redistributes RIP routes. You can define a single mapping for each route. • OSPF—Redistributes a different OSPF process. You can define a single mapping for each process. Select a process from the displayed list, then select one or more match criteria: – Internal—Routes that are internal to a specific AS. – External1—Routes that are external to the AS and imported into OSPF as a Type 1 external route. – External2—Routes that are external to the AS and imported into the selected process as a Type 2 external route. – NSAAExternal1—Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes. – NSAAExternal2—(NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes. • Connected—Redistributes routes that are established automatically by virtue of having enabled IP on an interface. These routes are redistributed as external to the AS.
Metric	A value representing the cost of the redistributed route. Valid values range from 0 to 4294967295.

EIGRP Routing on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Enhanced Interior Gateway Routing Protocol (EIGRP) is an enhanced distance vector protocol developed by Cisco Systems that integrates the capabilities of link-state protocols. EIGRP is suited for many different topologies and media. Key capabilities that distinguish EIGRP from other routing protocols are fast convergence, support for variable-length subnet masks, partial updates, and multiple network-layer protocols.

The metric that the router uses to reach the destination, and to advertise to other routers, is the sum of the best-advertised metrics from all neighbors and the link cost to the best neighbor.

EIGRP uses neighbor tables to store address and interface information about each of the router's neighbors. Hello packets advertise hold times, which is the length of time a neighbor can be considered reachable and operational. Topology tables contain all destinations advertised by neighboring routers. For each neighbor, the entry records the advertised metric, which the neighbor stores in its routing table.

A router running EIGRP stores all its neighbors' routing tables so that it can quickly adapt to alternate routes. If no appropriate route exists, EIGRP queries its neighbors to discover an alternate route. These queries propagate until an alternate route is found. EIGRP sends incremental updates when the state of a destination changes, instead of sending the entire contents of the routing table. EIGRP ensures that only those routers needing the information are updated. This feature minimizes the bandwidth required for EIGRP packets.

EIGRP supports both internal and external routes. Internal routes originate within an EIGRP Autonomous System (AS). Therefore, a directly attached network that is configured to run EIGRP is considered an internal route and is propagated with this information throughout the AS. External routes are learned by another routing protocol or reside in the routing table as static routes. These routes are tagged individually with the identity of their origin.

The following topics describe the tasks you perform to create an EIGRP routing policy:

- Defining EIGRP Routes, page 67-9
- Defining EIGRP Interface Properties, page 67-10
- Redistributing Routes into EIGRP, page 67-12

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- RIP Routing on Cisco IOS Routers, page 67-42
- OSPF Routing on Cisco IOS Routers, page 67-19
- BGP Routing on Cisco IOS Routers, page 67-1

Defining EIGRP Routes

To configure an EIGRP routing policy, you must assign each autonomous system a number, which identifies the autonomous system to other routers. You then must select the networks to which routes will be created. In addition, you can select which interfaces should be passive. Unlike other routing protocols, passive interfaces in EIGRP neither send nor receive routing updates from their neighbors, resulting in the loss of their neighbor relationship.

When you configure EIGRP routing policies, you can also decide whether to enable auto-summarization, which greatly simplifies routing tables and the exchange of routing information by having many subnets represented by a single network entry.

Related Topics

- Defining EIGRP Interface Properties, page 67-10
- Redistributing Routes into EIGRP, page 67-12
- EIGRP Routing on Cisco IOS Routers, page 67-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > EIGRP** from the Policy selector, then click the **Setup** tab in the work area.

- (Policy view) Select **Router Platform > Routing > EIGRP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The EIGRP Setup tab is displayed (see EIGRP Page—Setup Tab, page 67-13).

Step 2 On the EIGRP Setup tab, select an EIGRP route from the table, then click **Edit**, or click **Add** to create a route. The EIGRP Setup dialog box appears. See Table 67-6 on page 67-14 for a description of the fields in this dialog box.

Step 3 Enter the autonomous system number for the route. This number identifies the autonomous system to other routers.

Step 4 Enter the addresses of the networks to include in the EIGRP route. You can use a combination of addresses and network/host objects; separate addresses with commas. Click **Select** to select network/host objects from a list of existing objects, or to create new network/host objects. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 5 Enter the addresses of the passive interfaces, which are interfaces that should not send routing updates to their neighbors, if any. Enter the names of one or more interfaces or interface roles; separate addresses with commas. Click **Select** to select interface names or roles from a list of existing objects, or to create new interface role objects. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 6 (Optional) Select **Auto-Summary** to enable the auto summarization of subnet routes into network-level routes. Summarization reduces the size of routing tables, thereby reducing the complexity of the network.

Step 7 Click **OK** to save your definitions. The EIGRP route appears in the table displayed in the EIGRP Setup tab.

Defining EIGRP Interface Properties

You can optionally modify the default values of the following two interface properties in a selected EIGRP autonomous system:

- Hello interval.
- Split horizon.

The hello interval defines the interval between hello packets. Routing devices periodically send these packets to each other to dynamically learn of other routers on their directly attached networks. This information is used to discover neighbors and to learn when neighbors become unreachable or inoperative. By default, hello packets are sent every 5 seconds. The default interval for low speed (T1 or slower), nonbroadcast multiaccess (NBMA) media is every 60 seconds.

Split horizon is a feature that prevents route information from being sent back in the direction from which that information originated. If you enable split horizon on an interface (this is the default), update and query packets are not sent to destinations for which this interface is the next hop. This helps to prevent routing loops.

For example, as shown in Figure 67-1, if Router One is connected to Routers Two and Three through a single multipoint interface, and Router One learned about Network A from Router Two, Router One does not advertise the route to Network A over that same multipoint interface to Router Three. Router One assumes that Router Three would learn about Network A directly from Router Two.

Figure 67-1 EIGRP Split Horizon Example

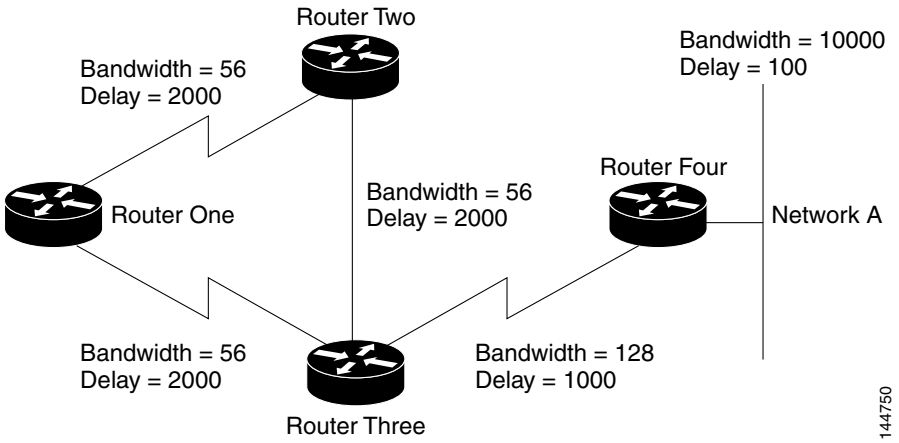

Split horizon is enabled by default on all EIGRP interfaces, because it typically optimizes communications among multiple routing devices. However, in certain cases with nonbroadcast networks (such as Frame Relay and SMDS), you might want to disable split horizon.

If you decide to disable split horizon on an EIGRP interface, keep the following in mind:

- In a hub-and-spoke network, you should disable split horizon only at the hub. This is because disabling split horizon on the spokes greatly increases EIGRP memory consumption on the hub router, as well as the amount of traffic generated on the spoke routers.
- Changing the split horizon setting on an interface resets all adjacencies with the EIGRP neighbors that are reachable over that interface.

Before You Begin

- Define at least one EIGRP autonomous system. See Defining EIGRP Routes, page 67-9.

Related Topics

- Defining EIGRP Routes, page 67-9
- Redistributing Routes into EIGRP, page 67-12
- EIGRP Routing on Cisco IOS Routers, page 67-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > EIGRP** from the Policy selector, then click the **Interfaces** tab in the work area.
- (Policy view) Select **Router Platform > Routing > EIGRP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Interfaces** tab.

The EIGRP Interfaces tab is displayed. See Table 67-7 on page 67-15 for a description of the fields on this tab.

Step 2 On the EIGRP Interfaces tab, select an interface from the table, then click **Edit**, or click **Add** to create an interface definition. The EIGRP Interface dialog box appears. See Table 67-8 on page 67-16 for a description of the fields in this dialog box.

Step 3 Select the AS number of the autonomous system whose interface properties you want to modify. See Defining EIGRP Routes, page 67-9 for more information about defining an autonomous system.

Step 4 Enter the name of the interface or interface role to define, or click **Select** to select an interface role from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 5 (Optional) In the Hello Interval field, modify the default interval between hello packets sent over the selected interfaces.

The default is 5 seconds for all interfaces, except for low-speed (T1 or less) NBMA media, for which the default interval is 60 seconds.

Step 6 (Optional) Deselect the **Split Horizon** check box to disable the split horizon feature. If you disable this feature, the selected interfaces can advertise a route out of the interface from which they learned the route.

Note In general, we recommend that you not disable split horizon unless you are certain that your application requires the change to properly advertise routes. If you disable split horizon on a serial interface, and that interface is attached to a packet-switched network, you must disable split horizon for all routers and access servers in all relevant multicast groups on that network.

Step 7 Click **OK** to save your definitions locally on the client and close the dialog box. The interface definition appears in the table on the EIGRP Interfaces tab.

Redistributing Routes into EIGRP

Redistribution refers to using a routing protocol, such as EIGRP, to advertise routes that are learned by some other means, such as a different routing protocol, static routes, or directly connected routes. For example, you can redistribute routes from the RIP routing protocol into your EIGRP autonomous system (AS). Redistribution is necessary in networks that operate in multiple-protocol environments and can be applied to all IP-based routing protocols.

Before You Begin

- Define at least one EIGRP autonomous system. See Defining EIGRP Routes, page 67-9.

Related Topics

- Defining EIGRP Routes, page 67-9
- Defining EIGRP Interface Properties, page 67-10
- EIGRP Routing on Cisco IOS Routers, page 67-8

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > EIGRP** from the Policy selector, then click the **Redistribution** tab in the work area.
- (Policy view) Select **Router Platform > Routing > EIGRP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Redistribution** tab.

The EIGRP Redistribution tab is displayed. See Table 67-9 on page 67-17 for a description of the fields on this tab.

Step 2 On the EIGRP Redistribution tab, select a row from the EIGRP Redistribution Mappings table, then click **Edit**, or click **Add** to create a mapping. The EIGRP Redistribution Mapping dialog box appears. See Table 67-10 on page 67-18 for a description of the fields in this dialog box.

Step 3 Select an existing EIGRP AS from the displayed list.

Step 4 Select the protocol whose routes you want to redistribute into the selected EIGRP AS.

Note You can create a single mapping for each static route, RIP route, BGP AS, EIGRP AS, and OSPF process.

Step 5 (Optional) Under Metrics, modify the default metric (cost) of the redistributed routes by entering values in the fields used to calculate the metric. The metric determines the priority of the routes.

Note Entering a metric is optional, but if you do specify a value, you must enter values for all five parameters. You need not define metric values when redistributing one EIGRP process into another.

Step 6 Click **OK** to save your definitions locally on the client and close the dialog box. The redistribution mapping appears in the Redistribution Mapping table in the EIGRP Redistribution tab.

EIGRP Routing Policy Page

Enhanced Interior Gateway Routing Protocol (EIGRP) is a scalable interior gateway protocol that provides extremely quick convergence times with minimal network traffic.

You can configure EIGRP routing policies from the following tabs on the EIGRP Routing page:

- EIGRP Page—Setup Tab, page 67-13
- EIGRP Page—Interfaces Tab, page 67-15
- EIGRP Page—Redistribution Tab, page 67-17

For more information, see EIGRP Routing on Cisco IOS Routers, page 67-8.

Navigation Path

- (Device view) Select **Platform > Routing > EIGRP** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > EIGRP** from the Policy Type selector. Right-click **EIGRP** to create a policy, or select an existing policy from the Shared Policy selector.

EIGRP Page—Setup Tab

Use the EIGRP Setup tab to view, create, edit, and delete EIGRP routes.

Navigation Path

Go to the EIGRP Routing Policy Page, page 67-13, then click the **Setup** tab.

Related Topics

- Defining EIGRP Routes, page 67-9
- EIGRP Page—Interfaces Tab, page 67-15
- EIGRP Page—Redistribution Tab, page 67-17
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-5 EIGRP Setup Tab

Element	Description
AS Number	The autonomous system number that identifies the autonomous system to other routers.
Networks	The names of the networks included in the route.
Passive Interfaces	The interfaces that neither send nor receive routing updates from their neighbors.
Auto-Summary	Indicates whether auto summarization is activated on the selected route.
Add button	Opens the EIGRP Setup Dialog Box, page 67-14. From here you can create an EIGRP route.
Edit button	Opens the EIGRP Setup Dialog Box, page 67-14. From here you can edit the selected EIGRP route.
Delete button	Deletes the selected EIGRP routes from the table.

EIGRP Setup Dialog Box

Use the EIGRP Setup dialog box to add or edit EIGRP routes.

Navigation Path

Go to the EIGRP Page—Setup Tab, page 67-13, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining EIGRP Routes, page 67-9
- Specifying IP Addresses During Policy Definition, page 6-87
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 67-6 EIGRP Setup Dialog Box

Element	Description
AS Number	The autonomous system number for the EIGRP route. This number is used to identify the autonomous system to other routers. Valid values are from 1 to 65535.

Table 67-6 EIGRP Setup Dialog Box (continued)

Element	Description
Networks	The networks associated with the EIGRP route. Enter one or more network addresses or network/host objects, separated by commas. Click **Select** to select network/host objects from a list of existing objects, or to create new objects.
Passive Interfaces	The interfaces that do not send updates to their routing neighbors. Enter one or more interface names or roles, separated by commas. Click **Select** to select interface names or roles from a list of existing objects, or to create new interface role objects. **Note** When you make an interface passive, EIGRP suppresses the exchange of hello packets between routers, resulting in the loss of their neighbor relationship. This not only stops routing updates from being advertised but also suppresses incoming routing updates.
Auto-Summary	When selected, enables the automatic summarization of subnet routes into network-level routes. Summarization reduces the size of routing tables, thereby reducing the complexity of the network. When deselected, automatic summarization is disabled.

EIGRP Page—Interfaces Tab

Use the EIGRP Interfaces tab to create, edit, and delete interface properties for selected EIGRP autonomous systems. This includes modifying the default hello interval and disabling split horizon.

Note You can access the EIGRP Interfaces tab only after defining at least one EIGRP autonomous system in the Setup tab. See EIGRP Page—Setup Tab, page 67-13.

Navigation Path

Go to the EIGRP Routing Policy Page, page 67-13, then click the **Interfaces** tab.

Related Topics

- Defining EIGRP Interface Properties, page 67-10
- EIGRP Page—Setup Tab, page 67-13
- EIGRP Page—Redistribution Tab, page 67-17
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-7 EIGRP Interfaces Tab

Element	Description
AS Number	The EIGRP autonomous system number for which interface properties are defined.

Table 67-7 EIGRP Interfaces Tab (continued)

Element	Description
Interfaces	The interfaces related to the selected EIGRP autonomous system that have specially defined values.
Split Horizon	Indicates whether the split horizon feature is enabled or disabled for the selected interface.
Hello Interval	The defined interval between hello packets sent to neighboring routers.
Add button	Opens the EIGRP Interface Dialog Box, page 67-16. From here you can create an EIGRP interface definition.
Edit button	Opens the EIGRP Interface Dialog Box, page 67-16. From here you can edit the selected EIGRP interface definition.
Delete button	Deletes the selected EIGRP interface definitions from the table.

EIGRP Interface Dialog Box

Use the EIGRP Interface dialog box to add or edit interface definitions for a selected EIGRP autonomous system.

Navigation Path

Go to the EIGRP Page—Interfaces Tab, page 67-15, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining EIGRP Interface Properties, page 67-10
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 67-8 EIGRP Interface Dialog Box

Element	Description
AS Number	Selects the EIGRP autonomous system number whose interface properties you want to modify. For more information about EIGRP autonomous systems, see EIGRP Setup Dialog Box, page 67-14.
Interface	Specifies the EIGRP interface you wish to configure. Enter the name of an interface or interface role, or click **Select** to select an interface role object from a list or to create a new one.
Hello Interval	The default interval between hello packets sent by the router to its neighbors. Routers send hello packets to each other to dynamically learn of other routers on their directly attached networks. Valid values range from 1 to 65535 seconds. The default is 5 seconds.

Table 67-8 EIGRP Interface Dialog Box (continued)

Element	Description
Split Horizon	When selected, the split horizon feature is used to prevent routing loops. When deselected, split horizon is disabled. When split horizon is disabled, the router can advertise a route out of the same interface through which it learned the route. Disabling split horizon is often useful when dealing with nonbroadcast networks, such as Frame Relay and SMDS. **Note** Changing the split horizon setting on an interface resets all adjacencies with EIGRP neighbors that are reachable over that interface.

EIGRP Page—Redistribution Tab

Use the EIGRP Redistribution tab to create, edit, and delete EIGRP redistribution mappings.

Navigation Path

Go to the EIGRP Routing Policy Page, page 67-13, then click the **Redistribution** tab.

Related Topics

- Redistributing Routes into EIGRP, page 67-12
- EIGRP Page—Setup Tab, page 67-13
- EIGRP Page—Interfaces Tab, page 67-15
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-9 EIGRP Redistribution Tab

Element	Description
EIGRP AS Number	The area ID of the EIGRP route into which other routes are being redistributed.
Protocol	The protocol that is being redistributed.
AS/Process ID	The AS number or process ID of the route being redistributed.
Bandwidth	The minimum bandwidth of the path for the EIGRP route, as defined for the route metric.
Delay	The mean latency of the path, as defined for the route metric.
Reliability	A value representing the estimated reliability of the path, as defined for the route metric.
Effective Bandwidth	A value representing the effective load on the link, as defined for the route metric.
MTU	The minimum MTU of the path, as defined for the route metric.

Table 67-9 EIGRP Redistribution Tab (continued)

Element	Description
Match	When redistributing an OSPF process, indicates the types of OSPF routes that are being redistributed.
Add button	Opens the EIGRP Redistribution Mapping Dialog Box, page 67-18. From here you can define EIGRP redistribution mappings.
Edit button	Opens the EIGRP Redistribution Mapping Dialog Box, page 67-18. From here you can edit the selected EIGRP redistribution mapping.
Delete button	Deletes the selected EIGRP redistribution mappings from the table.

EIGRP Redistribution Mapping Dialog Box

Use the EIGRP Redistribution Mapping dialog box to add or edit the properties of an EIGRP redistribution mapping.

Navigation Path

Go to the EIGRP Page—Redistribution Tab, page 67-17, then click the **Add** or **Edit** button beneath the table.

Note You must create at least one EIGRP AS before you can access the EIGRP Redistribution dialog box. See EIGRP Page—Setup Tab, page 67-13.

Related Topics

- Redistributing Routes into EIGRP, page 67-12

Field Reference

Table 67-10 EIGRP Redistribution Mapping Dialog Box

Element	Description
EIGRP AS Numbers	The EIGRP AS into which other routes are being redistributed. You must select an ID number from the list of EIGRP autonomous systems defined in the EIGRP Page—Setup Tab, page 67-13.
Protocol to Redistribute	The routing protocol that is being redistributed: • Static—Redistributes static routes. You can define a single mapping for each route. • EIGRP—Redistributes an EIGRP autonomous system. Enter the AS number in the displayed field. You can define a single mapping for each AS. • BGP—Redistributes a BGP autonomous system. You can define a single BGP mapping on each device. If you configured a BGP AS in the BGP Setup tab, the AS number is displayed. Otherwise, a message is displayed indicating that no BGP AS was defined. See BGP Page—Redistribution Tab, page 67-7.

Table 67-10 EIGRP Redistribution Mapping Dialog Box (continued)

Element	Description
Protocol to Redistribute (continued)	• OSPF—Redistributes a different OSPF process. You can define a single mapping for each process. Select a process from the displayed list, then select one or more match criteria: – Internal—Routes that are internal to a specific AS. – External1—Routes that are external to the AS and imported into OSPF as a Type 1 external route. – External2—Routes that are external to the AS and imported into the selected process as a Type 2 external route. – NSAAExternal1—Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes. – NSAAExternal2—(NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes. • RIP—Redistributes RIP routes. • Connected—Redistributes routes that are established automatically by virtue of having enabled IP on an interface. These routes are redistributed as external to the AS.
Metrics	The default metric (cost) of the redistributed route. Metric parameters include: • Bandwidth—The minimum bandwidth of the path in kilobits per second. Valid values range from 1 to 4294967295. • Delay—The mean latency of the path in units of 10 microseconds. Valid values range from 0 to 4294967295. • Reliability—A value expressing the estimated reliability of the link. Valid values range from 0 to 255, where 255 represents 100% reliability. • Effective Bandwidth—A value expressing the effective load on the link. Valid values range from 1 to 255, where 255 represents 100% utilization. • MTU of Path—The maximum transmission unit of the path. Valid values range from 1 to 65535 bytes.

OSPF Routing on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Open Shortest Path First (OSPF) is an interior gateway routing protocol that uses link states instead of distance vectors to distribute routing information within a single autonomous system (AS). OSPF propagates link-state advertisements (LSAs) instead of routing table updates, which allows OSPF networks to converge more quickly than RIP networks. You define areas to limit the number of LSAs that need to be propagated to changes that occur within the area.

A router that has interfaces in multiple OSPF areas is called an Area Border Router (ABR). An ABR uses LSAs to send information about available routes to other OSPF routers. A router that acts as a gateway to redistribute traffic between routers using OSPF and routers using other routing protocols is called an Autonomous System Boundary Router (ASBR). Any router can act as an ABR or ASBR.

The following topics describe the tasks you perform to create an OSPF routing policy:

- Defining OSPF Process Settings, page 67-20
- Defining OSPF Area Settings, page 67-21
- Redistributing Routes into OSPF, page 67-22
- Defining OSPF Interface Settings, page 67-25

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- RIP Routing on Cisco IOS Routers, page 67-42
- EIGRP Routing on Cisco IOS Routers, page 67-8
- BGP Routing on Cisco IOS Routers, page 67-1

Defining OSPF Process Settings

You configure OSPF process parameters by specifying a process ID number, which identifies the OSPF process to other routers, and by deciding whether any interfaces should be passive. Passive interfaces do not send routing updates to their neighbors.

Related Topics

- Defining OSPF Area Settings, page 67-21
- Defining OSPF Interface Settings, page 67-25
- Redistributing Routes into OSPF, page 67-22
- OSPF Routing on Cisco IOS Routers, page 67-19

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > OSPF Process** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Routing > OSPF Process** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The OSPF Process Setup tab is displayed. See Table 67-13 on page 67-35 for a description of the fields on this tab.

Step 2 On the OSPF Process Setup tab, select an OSPF process from the table, then click **Edit**, or click **Add** to create a process. The OSPF Setup dialog box appears. See Table 67-14 on page 67-36 for a description of the fields in this dialog box.

Step 3 Enter the process ID number in the field provided. The process ID defined here does not need to match the process ID on any other devices.

Step 4 Define which interfaces should not send routing updates to its neighbors:

a. Click **Edit** under Passive Interfaces to display the Edit Interfaces dialog box. Use this dialog box to define which interfaces should *not* send routing updates to its neighbors.

b. Enter the names of one or more interfaces or interface roles, or click **Select** to select an interface role from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

c. Click **OK** to save your changes and return to the OSPF Setup dialog box.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box.

Defining OSPF Area Settings

You configure OSPF area settings by associating an area ID with a particular OSPF process, selecting the networks included in the area, and selecting the type of authentication used by the routers in the area.

Each OSPF process that you define should contain at least one defined area. If you define more than one area, one area must be area 0. This is called the backbone. All other areas must be physically connected to the backbone. This enables other areas to inject routing information into the backbone, which the backbone distributes to the remaining areas.

You must configure at least one OSPF process before defining OSPF area/network settings for that process.

Related Topics

- Defining OSPF Process Settings, page 67-20
- Defining OSPF Interface Settings, page 67-25
- Redistributing Routes into OSPF, page 67-22
- OSPF Routing on Cisco IOS Routers, page 67-19

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > OSPF Process** from the Policy selector, then click the **Area** tab in the work area.
- (Policy view) Select **Router Platform > Routing > OSPF Process** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Area** tab.

The OSPF Process Area tab is displayed. See Table 67-15 on page 67-37 for a description of the fields on this tab.

Step 2 On the OSPF Process Area tab, select an OSPF area from the table, then click **Edit**, or click **Add** to create an area. The OSPF Area dialog box appears. See Table 67-16 on page 67-37 for a description of the fields in this dialog box.

Step 3 Select a process ID from the displayed list.

Step 4 Enter an area ID to associate with the selected OSPF process.

Step 5 Enter the addresses of the networks to include in the OSPF area. You can enter a combination of addresses and network/host objects, or click **Select** to select a network/host object from a list or to create a new one. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 6 Select the authentication type to use in the OSPF area: MD5, clear text, or none. We recommend MD5 when security is of concern. Please note the following:

- The authentication type must be the same for all routers and access servers in the same area.
- Specifying clear-text authentication for an area sets the authentication to Type 1 (simple password). All routers on a network must use the same clear-text password to communicate with each other using OSPF.
- MD5 passwords need not be the same throughout an area, but they must be the same between neighbors.
- If you use interface authentication (see Defining OSPF Interface Settings, page 67-25), the authentication type used for the area must match the authentication type used for the interface.

Step 7 Click **OK** to save your definitions. The OSPF area appears in the table displayed on the OSPF Area tab.

Redistributing Routes into OSPF

Redistribution refers to using a routing protocol, such as OSPF, to advertise routes that are learned by some other means, such as a different routing protocol, static routes, or directly connected routes. For example, you can redistribute routes from the RIP routing protocol into your OSPF domain. Redistribution is necessary in networks that operate in multiple-protocol environments and can be applied to all IP-based routing protocols.

Redistributing routes into OSPF from other routing protocols or from static routes causes these routes to become OSPF external routes (Type 1 or Type 2).

Redistributing routes into OSPF involves:

- Defining OSPF Redistribution Mappings, page 67-22
- Defining OSPF Maximum Prefix Values, page 67-24

Related Topics

- Defining OSPF Process Settings, page 67-20
- Defining OSPF Area Settings, page 67-21
- Defining OSPF Interface Settings, page 67-25
- OSPF Routing on Cisco IOS Routers, page 67-19

Defining OSPF Redistribution Mappings

When you define OSPF redistribution mappings, you must select the protocol to redistribute and the OSPF process into which routes from that protocol are redistributed. Additionally, you can manually define the metric, which determines the priority of the redistributed routes, and the type of external OSPF route to create, Type 1 or Type 2.

You can create multiple mappings to the same OSPF process. For example, you can redistribute both RIP and EIGRP routes into the same OSPF process. You can also redistribute routes from other OSPF processes.

Note Redistribution into an OSPF Not-So-Stubby Area (NSSA) creates a special type of link-state advertisement (LSA) called type 7, which can exist only in an NSSA area. An NSSA autonomous system router (ASBR) generates this LSA, and an NSSA area border router (ABR) translates it into a type 5 LSA, which is propagated into the OSPF domain.

Type 1 versus Type 2 External Routes

Two types of OSPF external routes exist, Type 1 and Type 2. The difference between the two is related to how the cost (metric) of the route is calculated. The cost of a Type 1 route is the sum of the external cost and the internal cost used to reach that route. The cost of a Type 2 route is based on the external cost only. By default, external routes are defined as Type 2. However, a Type 1 route is always preferred over a Type 2 route to the same destination.

Before You Begin

- Define at least one OSPF process. See Defining OSPF Process Settings, page 67-20.

Related Topics

- Defining OSPF Maximum Prefix Values, page 67-24
- Redistributing Routes into OSPF, page 67-22

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > OSPF Process** from the Policy selector, then click the **Redistribution** tab in the work area.
- (Policy view) Select **Router Platform > Routing > OSPF Process** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Redistribution** tab.

The OSPF Process Redistribution tab is displayed. See Table 67-17 on page 67-38 for a description of the fields on this tab.

Step 2 On the OSPF Process Redistribution tab, select a row from the OSPF Redistribution Mappings table, then click **Edit**, or click **Add** to create a mapping. The OSPF Redistribution Mapping dialog box is displayed. See Table 67-18 on page 67-40 for a description of the fields in this dialog box.

Step 3 Select an existing OSPF process from the displayed list.

Step 4 Select the protocol whose routes you want to redistribute into the selected OSPF process.

Note You can create a single mapping for each static route, RIP route, BGP AS, EIGRP AS, and OSPF process.

Step 5 (Optional) Modify the default metric (cost) of the redistributed routes. The metric determines the priority of the routes.

Step 6 Select the Metric Type of external route to create, Type 1 or Type 2. The default is Type 2.

Step 7 (Optional) Select the **Limit to Subnets** check box to redistribute only subnetted routes. By default, this option is not selected.

Step 8 Click **OK** to save your definitions. The redistribution mapping appears in the Redistribution Mapping table on the OSPF Process Redistribution tab.

Defining OSPF Maximum Prefix Values

You can define a maximum number of prefixes (routes) that may be redistributed from other protocols or OSPF processes into a selected OSPF process. Setting a limit helps prevent the router from being flooded by too many redistributed routes. For example, without a defined maximum, flooding can occur when BGP is redistributed into OSPF.

When you define a maximum prefix value, you can decide whether to prevent additional routes from being redistributed once this maximum is reached, or whether to only issue a warning.

The redistribution limit applies to all IP redistributed prefixes, including summarized ones. The limit does not apply to default routes or prefixes that are generated as a result of type 7 to type 5 translations.

Before You Begin

- Define at least one OSPF process. Define at least one OSPF process. See Defining OSPF Process Settings, page 67-20.
- Define at least one OSPF redistribution mapping. See Defining OSPF Redistribution Mappings, page 67-22.

Related Topics

- Defining OSPF Redistribution Mappings, page 67-22
- Redistributing Routes into OSPF, page 67-22

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > OSPF Process** from the Policy selector, then click the **Redistribution** tab in the work area.
- (Policy view) Select **Router Platform > Routing > OSPF Process** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Redistribution** tab.

The OSPF Process Redistribution tab is displayed. See Table 67-17 on page 67-38 for a description of the fields on this tab.

Step 2 On the OSPF Process Redistribution tab, select a row from the Max Prefix Mapping table, then click **Edit**, or click **Add** to create a definition. The Max Prefix Mapping dialog box appears. See Table 67-19 on page 67-41 for a description of the fields in this dialog box.

Step 3 Select an existing OSPF process from the displayed list.

Step 4 In the Max Prefix field, enter the maximum number of routes that can be redistributed into the selected OSPF process.

Step 5 (Optional) Modify the default threshold percentage. When the number of redistributed routes reaches this threshold, a warning is issued. By default, the threshold value is 75% of the defined maximum prefix value.

Step 6 (Optional) Select what should happen when the maximum prefix value is reached:

- Enforce Maximum Route—Prevents additional routes from being redistributed to the selected process.
- Warning Only—Issues an additional warning, but allows route redistribution to continue even after the maximum prefix value is reached.

Note Flooding can result if you allow route redistribution to continue after exceeding the maximum prefix value.

Step 7 Click **OK** to save your definitions. The maximum prefix definition appears in the Maximum Prefix table on the OSPF Process Redistribution tab.

Defining OSPF Interface Settings

You can modify a variety of interface-specific OSPF parameters. This procedure describes how to define these parameters. For more information about a particular parameter, see the following topics:

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Understanding OSPF Timer Settings, page 67-28
- Blocking LSA Flooding, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29

Related Topics

- Defining OSPF Process Settings, page 67-20
- Defining OSPF Area Settings, page 67-21
- Redistributing Routes into OSPF, page 67-22
- OSPF Routing on Cisco IOS Routers, page 67-19

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > OSPF Interface** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > OSPF Interface** from the Policy Type selector. Select an existing policy or create a new one.

The OSPF Interface page is displayed. See Table 67-11 on page 67-30 for a description of the fields on this page.

Step 2 On the OSPF Interface page, select an interface definition from the table, then click **Edit**, or click **Add** to create a definition. The OSPF Interface dialog box appears. See Table 67-12 on page 67-32 for a description of the fields in this dialog box.

Step 3 Enter the name of the interface or interface role to define, or click **Select** to select an interface role from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 Define interface authentication. The authentication type you select for the interface must match the authentication type you select for the area (see Defining OSPF Area Settings, page 67-21).

All neighboring routers on the same network must have the same password to be able to exchange OSPF information. For more information, see Understanding OSPF Interface Authentication, page 67-29.

The key ID number can be associated with multiple passwords. This is an easy and secure way to migrate passwords. For example, to migrate from one password to another, configure a password under a different key ID, then remove the first key.

Tip The key ID number can be associated with multiple passwords. This is an easy and secure way to migrate passwords. For example, to migrate from one password to another, configure a password under a different key ID, then remove the first key.

Note Do not use clear text authentication in OSPF packets for security purposes, because the unencrypted authentication key is sent in every packet. Use clear text authentication only when security is not an issue, for example, to ensure that misconfigured hosts do not participate in routing.

Step 5 (Optional) Under **Properties**, configure interface parameters as required. See Table 67-12 on page 67-32 for information about each parameter.

Step 6 Click **OK** to save your definitions. The defined interfaces appear on the OSPF Interface page.

Step 7 Repeat the process to define interface-specific parameters on additional OSPF interfaces.

Understanding Interface Cost

The cost of an OSPF interface is a metric representing the cost of sending a packet over that interface. By default, this cost is calculated using this formula:

10^8 / bandwidth [bits per second]

For example, if the bandwidth of a Fast Ethernet interface is 10 Mbps (equal to 10^7), the cost of sending packets over that interface is calculated as $10^8 / 10^7$ or 10. This formula establishes an inverse relationship between the bandwidth of an interface and its cost; the greater the bandwidth, the lower the cost.

Although cost is a calculated value, you can manually enter the cost of a selected interface.

Related Topics

- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Blocking LSA Flooding, page 67-28
- Understanding OSPF Timer Settings, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Understanding Interface Priority

Routers that share a common segment are elected through the Hello protocol to be neighbors on that segment. Election occurs as soon as the routers see themselves listed in their neighbor's hello packet. Adjacency is the next step. Adjacent routers are routers that proceed beyond the simple Hello exchange to a database exchange.

On each multiaccess (as opposed to point-to-point) segment, OSPF elects one router as the designated router (DR) for that segment. The DR acts as a central point of contact to minimize information exchange. Each router in the segment sends updates to the DR, which in turn relays the information to the other routers. A second router is elected as the backup designated router (BDR) in case the DR goes down.

DR and BDR election is performed via the Hello protocol. The router with the highest OSPF priority becomes the DR for that segment. The same process is then repeated for the BDR. In the case of a tie, the router with the higher router ID (RID) is elected. By default, each interface is given a priority of 1, but you can assign a higher priority to selected interfaces, as required.

Note The priority setting does not apply to point-to-point, nonbroadcast interfaces.

Related Topics

- Understanding Interface Cost, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Blocking LSA Flooding, page 67-28
- Understanding OSPF Timer Settings, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Disabling MTU Mismatch Detection

The MTU is the largest packet size that a particular interface can handle. If one router sends a DBD packet that is larger than the MTU setting on a neighboring router, the neighboring router ignores the packet. In many cases, an MTU mismatch causes the two routers to become stuck in exstart/exchange state, which prevents OSPF adjacency from being established. This is why it is important that all neighboring routers share the same MTU setting and that MTU mismatch detection be enabled.

You can, however, disable MTU mismatch detection. This is useful in cases where mismatch detection is preventing adjacency from taking place in an otherwise valid setup between two devices with different MTUs.

Related Topics

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Blocking LSA Flooding, page 67-28
- Understanding OSPF Timer Settings, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Blocking LSA Flooding

By default, OSPF floods new LSAs over all interfaces in the same area, except the interface on which the LSA arrives. Although some redundancy is desirable, too much redundancy can waste bandwidth. In certain topologies, such as full mesh, LSA flooding can destabilize the network because of excessive link and CPU usage. Therefore, you can block LSA flooding to selected interfaces on broadcast, nonbroadcast, and point-to-point networks.

Related Topics

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Understanding OSPF Timer Settings, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Understanding OSPF Timer Settings

OSPF uses a series of timers during operation:

- Hello Interval—Determines how often an interface sends hello packets, which are used to acquire neighbors and act as indicators that the router is still functioning. The smaller the interval, the faster topological changes on the network are detected. However, a smaller interval also results in more traffic being sent over the interface. The hello interval must be the same on all routers and access servers on a specific network.
- Transmit Delay—Determines the delay before an LSA is flooded over the link. The transmit delay setting should take into account the transmission and propagation delays for the interface. These factors are particularly important when configuring low-speed and on-demand links.
- Retransmit Interval—Determines how long to wait before retransmitting an unacknowledged database description (DBD) packet to its neighbors. The retransmit interval setting should be low enough to prevent excessive retransmissions.

  Note You should increase the retransmit interval for serial lines and virtual links.

- Dead Interval—Determines how long an interface should wait before declaring its neighbor to be down. This declaration is caused by an absence of hello packets from the neighbor during this interval. The dead interval setting must be the same for all routers and access servers on a specific network. By default, this interval is four times the hello interval.

Related Topics

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Blocking LSA Flooding, page 67-28
- Understanding the OSPF Network Type, page 67-29

- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Understanding the OSPF Network Type

You can manually configure the OSPF network type on an interface as either broadcast or nonbroadcast multiaccess (NBMA), regardless of the default media type. For example, you can use this feature to configure broadcast networks (such as Ethernet, Token Ring, and FDDI) as NBMA when your network contains routers that do not support multicast addressing. You can also configure NBMA networks (such as X.25, Frame Relay, and SMDS) as broadcast networks, which eliminates the need to configure neighbors.

Configuring NBMA networks as either broadcast or nonbroadcast assumes the existence of virtual circuits (VCs) from every router to every router (fully meshed network). If VCs do not exist between each router, due to cost constraints or the existence of an only partially meshed network, you can configure the OSPF network type as point-to-multipoint. An OSPF point-to-multipoint interface is defined as a numbered point-to-point interface having one or more neighbors. It creates multiple host routes.

If you use the point-to-multipoint network type, routing between two routers that are not directly connected go through a third router that has VCs to both routers. You do not need to configure neighbors when using this feature. OSPF point-to-multipoint networks have the following benefits compared to NBMA and point-to-point networks:

- Point-to-multipoint is easier to configure because it consumes only one IP subnet and does not require neighbor configuration or designated router election.
- It costs less because it does not require a fully meshed topology.
- It is more reliable because it maintains connectivity in the event of VC failure.

Note For point-to-multipoint, broadcast networks, you can optionally define neighbors, in which case you should specify the cost to each neighbor. For point-to-multipoint, nonbroadcast networks, you must identify neighbors, but specifying a cost to each neighbor is optional. In both cases, you define neighbors using FlexConfig. See Chapter 7, "Managing FlexConfigs" for more information.

Related Topics

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Blocking LSA Flooding, page 67-28
- Understanding OSPF Timer Settings, page 67-28
- Understanding OSPF Interface Authentication, page 67-29
- Defining OSPF Interface Settings, page 67-25

Understanding OSPF Interface Authentication

You define neighbor authentication settings for OSPF interfaces by selecting the interfaces and selecting an authentication type, either MD5 or clear text.

When you use MD5 authentication, neighboring routers must share the same password. When you use clear-text authentication, all routers on the network using OSPF must share the same password.

Whenever you configure an interface with a new key, the router sends multiple copies of the same packet, each authenticated by different keys. The router stops sending duplicate packets when it detects that all of its neighbors have adopted the new key.

Note You should use authentication with all routing protocols when possible, because attackers can use route redistribution between OSPF and other protocols (such as RIP) to subvert routing information.

Related Topics

- Understanding Interface Cost, page 67-26
- Understanding Interface Priority, page 67-26
- Disabling MTU Mismatch Detection, page 67-27
- Blocking LSA Flooding, page 67-28
- Understanding OSPF Timer Settings, page 67-28
- Understanding the OSPF Network Type, page 67-29
- Understanding OSPF Interface Authentication, page 67-29

OSPF Interface Policy Page

Use the OSPF Interface page to view, create, edit, and delete interface-specific OSPF settings. For more information, see Defining OSPF Interface Settings, page 67-25.

Navigation Path

- (Device view) Select **Platform > Routing > OSPF Interface** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > OSPF Interface** from the Policy Type selector. Right-click **OSPF Interface** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- OSPF Process Policy Page, page 67-34
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-11 OSPF Interface Page

Element	Description
Interfaces	The name of an interface (as defined by an interface role) on which OSPF is enabled.
Authentication	The type of OSPF neighbor authentication enabled for the selected interface.

Table 67-11 OSPF Interface Page (continued)

Element	Description
Key ID	The identification number of the authentication key used for MD5 authentication.
Cost	The cost of sending packets over the selected interface, if this value is different from the cost as normally calculated.
Priority	The priority of the selected interface.
MTU Ignore	Indicates whether Maximum Transmission Rate (MTU) detection is disabled on the selected interface.
Database Filter	Indicates whether link-state advertisement (LSA) flooding is disabled on the selected interface.
Hello Interval	The interval between hello packets (in seconds) sent over this interface.
Transmit Delay	The amount of time OSPF waits (in seconds) before flooding an LSA over the link.
Retransmit Interval	The interval between LSA retransmissions (in seconds) over the selected interface.
Dead Interval	The interval OSPF waits (in seconds) before declaring a neighboring router dead because of an absence of hello packets.
Network Type	The network type configured for the selected interface, if it differs from the default medium.
Add button	Opens the OSPF Interface Dialog Box, page 67-31. From here you can define the properties of an OSPF interface.
Edit button	Opens the OSPF Interface Dialog Box, page 67-31. From here you can edit the properties of the selected OSPF interface.
Delete button	Deletes the selected OSPF interface definitions from the table.

OSPF Interface Dialog Box

Use the OSPF Interface dialog box to add or edit the properties of OSPF interfaces.

Navigation Path

Go to the OSPF Interface Policy Page, page 67-30, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining OSPF Interface Settings, page 67-25
- OSPF Routing on Cisco IOS Routers, page 67-19
- Basic Interface Settings on Cisco IOS Routers, page 62-1
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 67-12 OSPF Interface Dialog Box

Element	Description
Interface	The OSPF interface to configure. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one.
Authentication	Type—The authentication type used by the selected interface: • MD5—Uses the MD5 hash algorithm for authentication. This is the default. • Clear Text—Uses a clear text password for authentication. • None—Uses no authentication. **Note** The authentication type used on an interface must match the authentication type defined for the area. **Note** Use plain text authentication only when security is not an issue, for example, to ensure that misconfigured hosts do not participate in routing. • Key ID—Available only when MD5 is selected as the authentication type. The identification number of the authentication key. This number must be shared with all other devices sending updates to, and receiving updates from, the selected device. Valid values range from 1 to 255. • Key—The shared key used for authentication (MD5 or clear text). This key must be shared with all other devices sending updates to, and receiving updates from, the selected device. Enter this key again in the Confirm field. When using clear text, the key can include any continuous string of characters that can be entered from the keyboard (up to 8 bytes). When using MD5, the key can include alphanumeric characters only (up to 16 bytes).
Cost	The cost of sending packets over this interface. A value entered here overrides the default calculated cost (10 8 /bandwidth in bits per second). Valid values range from 1 to 65535.
Priority	The default priority of the interface. The priority is used to determine which routers become the designated router (DR) and backup designated router (BDR) for that segment. The higher the number, the higher the priority. The default priority is 1. Valid values range from 0 to 255. **Note** To exclude the interface from election as DR or BDR, assign a priority of 0. Configure router priority only for interfaces to multiaccess networks, not point-to-point networks.

Table 67-12 OSPF Interface Dialog Box (continued)

Element	Description
MTU Ignore	When selected, ignores MTU mismatches between neighboring routers. When deselected, MTU mismatch detection is enabled. **Note** Typically, this option is not used, because it can cause routers to become stuck in exstart/exchange state, which prevents OSPF adjacency from being established.
Database Filter	When selected, blocks link-state advertisement (LSA) flooding to the selected interface. When deselected, LSA flooding is permitted. **Note** We recommend that you enable this option on fully-meshed networks. This option is not available for point-to-multipoint networks.
Hello Interval	The default interval (in seconds) between hello packets sent over the selected interface. These packets are used by neighboring routers to confirm the router sending the packets is still operating. Valid values range 1 to 65535 seconds. **Note** The hello interval must be the same for all routers and access servers in the network.
Transmit Delay	The amount of time OSPF waits (in seconds) before flooding an LSA over the link. The default is 1 second. Valid values range from 1 to 65535 seconds. **Note** When you configure slow links or on-demand links that queue traffic before sending it in bursts, we recommend that you take these link delays into account when defining this value.
Retransmit Interval	The interval between LSA retransmissions (in seconds) over the selected interface. The default is 5 seconds. Valid values range from 1 to 65535 seconds. **Note** We recommend that you increase this value for serial lines and virtual links.
Dead Interval	The interval (in seconds) after which an interface declares its neighbor dead if no hello packets are received. Valid values range from 1 to 655335 seconds. **Note** The value of the dead interval is typically the hello interval value multiplied by 4. The dead interval must be the same for all routers and access servers in the network.

Table 67-12 OSPF Interface Dialog Box (continued)

Element	Description
Configure Network Type	When selected, enables you to select a network type that differs from the default medium used by the interface. When deselected, the network type is equivalent to the default medium used by the interface. For nonbroadcast multiaccess (NBMA) networks (such as ATM and Frame Relay), options are: • Broadcast—Treats the NBMA network as a broadcast network, which eliminates the need to configure neighbors. Use this option when there are virtual circuits from every router to every router (fully meshed network). • Point-to-Multipoint—Treats the nonbroadcast network as a series of point-to-point links. This option is easier to configure, less costly, and more reliable than NBMA or point-to-point networks. • Point-to-Multipoint Non-Broadcast—Statically maintains the known neighbors of the network. Selecting this option helps avoid the problem of losing neighbors that were learned dynamically through the reception of hello packets. **Note** Another option for NBMA networks is to configure neighbors manually using FlexConfigs. See Understanding FlexConfig Policies and Policy Objects, page 7-2. For broadcast networks (such as Ethernet, Token Ring, and FDDI), you can select: • Non-Broadcast—Treats the broadcast network as a nonbroadcast network. • Point-to-Point—Treats the broadcast network as a point-to-point network. You can use this option, for example, to configure a broadcast network (such as Ethernet) as a nonbroadcast multiaccess (NBMA) network if not all routers in the network support multicast addressing.

OSPF Process Policy Page

OSPF is an interior gateway routing protocol that uses link states instead of distance vectors for path selection. OSPF propagates link-state advertisements (LSAs) instead of routing table updates, which enables OSPF networks to converge quickly.

You can configure OSPF process policies from the following tabs on the OSPF Process page:

- OSPF Process Page—Setup Tab, page 67-35
- OSPF Process Page—Area Tab, page 67-36
- OSPF Process Page—Redistribution Tab, page 67-38

For more information, see OSPF Routing on Cisco IOS Routers, page 67-19.

Note For more information about OSPF interface policies, see OSPF Interface Policy Page, page 67-30.

Navigation Path

- (Device view) Select **Platform > Routing > OSPF Process** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > OSPF Process** from the Policy Type selector. Right-click **OSPF Process** to create a policy, or select an existing policy from the Shared Policy selector.

OSPF Process Page—Setup Tab

Use the OSPF Process Setup tab to create, edit, and delete OSPF processes. This includes selecting those interfaces that will remain passive, which means that they will not send routing updates to their neighbors. You can create as many processes for each router as required.

Navigation Path

Go to the OSPF Process Policy Page, page 67-34, then click the **Setup** tab.

Related Topics

- Defining OSPF Process Settings, page 67-20
- OSPF Process Page—Area Tab, page 67-36
- OSPF Process Page—Redistribution Tab, page 67-38
- OSPF Interface Policy Page, page 67-30
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-13 OSPF Process Setup Tab

Element	Description
Process ID	The process ID that identifies the OSPF routing process to other routers.
Passive Interfaces	The interfaces that do not send out routing updates.
Add button	Opens the OSPF Setup Dialog Box, page 67-35. From here you can define an OSPF process.
Edit button	Opens the OSPF Setup Dialog Box, page 67-35. From here you can edit the selected OSPF process.
Delete button	Deletes the selected OSPF processes from the table.

OSPF Setup Dialog Box

Use the OSPF Setup dialog box to add or edit an OSPF process.

Navigation Path

Go to the OSPF Process Page—Setup Tab, page 67-35, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining OSPF Process Settings, page 67-20

Field Reference

Table 67-14 OSPF Setup Dialog Box

Element	Description
Process ID	The process ID number for the OSPF process. This number identifies the OSPF process to other routers. It does not need to match the process ID on other devices. Valid values are from 1 to 65535.
Passive Interfaces	The interfaces that do not send updates to their routing neighbors. Click **Edit** to display the Edit Interfaces Dialog Box—OSPF Passive Interfaces, page 67-36. From here you can define these interfaces. **Note** When you make an interface passive, OSPF suppresses the sending of hello packets to neighboring routers. The interface will continue to receive routing updates, however.

Edit Interfaces Dialog Box—OSPF Passive Interfaces

When you configure an OSPF routing policy on a Cisco IOS router, use the Edit Interfaces dialog box to specify which interfaces will not send updates to their routing neighbors. Separate multiple names or roles with commas. Click **Select** to select interface names or roles from a list of existing objects, or to create new interface role objects.

Navigation Path

Go to the OSPF Setup Dialog Box, page 67-35, then click the **Edit** button in the Passive Interfaces field.

Related Topics

- OSPF Process Page—Setup Tab, page 67-35
- Defining OSPF Process Settings, page 67-20

OSPF Process Page—Area Tab

Use the OSPF Area tab to create, edit, and delete the areas and networks contained in each OSPF process. This includes selecting the type of authentication used by each area.

Navigation Path

Go to the OSPF Process Policy Page, page 67-34, then click the **Area** tab.

Related Topics

- Defining OSPF Area Settings, page 67-21
- OSPF Process Page—Setup Tab, page 67-35

- OSPF Process Page—Redistribution Tab, page 67-38
- OSPF Interface Policy Page, page 67-30
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-15 OSPF Process Area Tab

Element	Description
Area ID	The ID number of the area associated with the process.
Process ID	The process ID that identifies the OSPF routing process to other routers.
Networks	The networks included in the area.
Authentication	The authentication type used by the area—MD5, clear text, or none.
Add button	Open the OSPF Area Dialog Box, page 67-37. From here you can define an OSPF area.
Edit button	Opens the OSPF Area Dialog Box, page 67-37. From here you can edit the selected OSPF area.
Delete button	Deletes the selected OSPF areas from the table.

OSPF Area Dialog Box

Use the OSPF Area dialog box to add or edit the properties of an OSPF area. You should define at least one area for each OSPF process (see OSPF Setup Dialog Box, page 67-35), but deployment will not fail if you do not.

Navigation Path

Go to the OSPF Process Page—Area Tab, page 67-36, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining OSPF Area Settings, page 67-21
- Specifying IP Addresses During Policy Definition, page 6-87
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 67-16 OSPF Area Dialog Box

Element	Description
Process ID	The process ID associated with the OSPF area. The list contains the OSPF processes defined in the OSPF Process Page—Setup Tab, page 67-35.
Area ID	The area ID number associated with the selected process. Valid values range from 0 to 4294967295.

Table 67-16 OSPF Area Dialog Box (continued)

Element	Description
Networks	The networks to add to the OSPF area. Enter one or more network addresses or network/host objects, or click **Select** to select the object from a list or to create a new one.
Authentication	The type of authentication used for the area: • MD5—(Recommended) Uses the MD5 hash algorithm for authentication. • Clear Text—Uses clear text for authentication. • None—No authentication is used. **Note** The authentication type must be the same for all routers and access servers in an area.

OSPF Process Page—Redistribution Tab

Use the OSPF Process Redistribution tab to create, edit, and delete OSPF redistribution mappings. This includes defining the maximum number of routes that can be redistributed into OSPF from other protocols or other OSPF processes.

Navigation Path

Go to the OSPF Process Policy Page, page 67-34, then click the **Redistribution** tab.

Related Topics

- Redistributing Routes into OSPF, page 67-22
- OSPF Process Page—Setup Tab, page 67-35
- OSPF Process Page—Area Tab, page 67-36
- OSPF Interface Policy Page, page 67-30
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-17 OSPF Process Redistribution Tab

Element	Description
OSPF Redistribution Mapping Table	
OSPF Process ID	The ID of the OSPF routing domain into which other routes are being redistributed.
Protocol	The protocol that is being redistributed.
AS/Process ID	The AS number or process ID of the route that is being redistributed.
Match	When redistributing an OSPF process, indicates the types of OSPF routes that are being redistributed.
Metric	The value that determines the priority of the redistributed route.

Table 67-17 OSPF Process Redistribution Tab (continued)

Element	Description
Metric Type	The external link type associated with the default route advertised into the OSPF routing domain.
Subnets	Indicates whether routes that are subnetted are also being redistributed.
Add button	Opens the OSPF Redistribution Mapping Dialog Box, page 67-39. From here you can define OSPF redistribution mappings.
Edit button	Opens the OSPF Redistribution Mapping Dialog Box, page 67-39. From here you can edit the selected OSPF redistribution mapping.
Delete button	Deletes the selected redistribution mappings from the table.
OSPF Max Prefix Mapping Table	
OSPF Process ID	The ID of the OSPF routing domain for which a maximum prefix values has been defined.
Max Prefix	The maximum number of prefixes (routes) that may be redistributed to the selected OSPF process.
Threshold	The percentage of the maximum prefix value that acts as a threshold for triggering a warning message.
Action	Indicates whether redistribution to this OSPF process will stop when the maximum is reached, or whether only a warning is displayed.
Add button	Opens the OSPF Max Prefix Mapping Dialog Box, page 67-41. From here you can define maximum prefix values for OSPF processes.
Edit button	Opens the OSPF Max Prefix Mapping Dialog Box, page 67-41. From here you can edit the maximum prefix value defined for the selected OSPF process.
Delete button	Deletes the selected max prefix mappings from the table.

OSPF Redistribution Mapping Dialog Box

Use the OSPF Redistribution Mapping dialog box to add or edit the properties of an OSPF redistribution mapping.

Navigation Path

Go to the OSPF Process Page—Redistribution Tab, page 67-38, then click the **Add** or **Edit** button beneath the Redistribution Mapping table.

Note You must create at least one OSPF process before you can access the OSPF Redistribution dialog box. See OSPF Process Page—Setup Tab, page 67-35.

Related Topics

- OSPF Max Prefix Mapping Dialog Box, page 67-41
- Redistributing Routes into OSPF, page 67-22

Field Reference

Table 67-18 OSPF Redistribution Mapping Dialog Box

Element	Description
Process ID	The OSPF process into which other routes are being redistributed. You must select a process ID number from the list of OSPF processes defined in the OSPF Process Page—Setup Tab, page 67-35.
Protocol to Redistribute	The routing protocol that is being redistributed: • Static—Redistributes static routes. You can define a single mapping for each route. • EIGRP—Redistributes an EIGRP autonomous system. Enter the AS number in the displayed field. You can define a single mapping for each AS. • BGP—Redistributes a BGP autonomous system. You can define a single BGP mapping on each device. If you configured a BGP AS in the BGP Setup tab, the AS number is displayed. Otherwise, a message is displayed indicating that no BGP AS was defined. See BGP Page—Redistribution Tab, page 67-7.
Protocol to Redistribute (continued)	• OSPF—Redistributes a different OSPF process. You can define a single mapping for each process. Select a process from the displayed list, then select one or more match criteria: – Internal—Routes that are internal to a specific AS. – External1—Routes that are external to the AS and imported into OSPF as a Type 1 external route. – External2—Routes that are external to the AS and imported into the selected process as a Type 2 external route. – NSAAExternal1—Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes. – NSAAExternal2—(NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes. • RIP—Redistributes RIP routes. You can define a single mapping for each route. • Connected—Redistributes routes that are established automatically by virtue of having enabled IP on an interface. These routes are redistributed as external to the AS.
Default Metric	A value representing the cost of the redistributed route.
Metric Type	The external link type that is associated with the route being redistributed into the OSPF routing domain: • 1—Type 1 external route. The metric is the sum of the external redistributed cost and the internal OSPF cost. • 2—Type 2 external route. The metric is equal to the external redistributed cost, as defined in the Metric field. This is the default.

Table 67-18 OSPF Redistribution Mapping Dialog Box (continued)

Element	Description
Limit to Subnets	When selected, only subnetted routes are redistributed. When deselected, subnetted routes are not redistributed.

OSPF Max Prefix Mapping Dialog Box

Use the OSPF Max Prefix Mapping dialog box to add or edit the maximum number of routes that can be redistributed into an OSPF process.

Navigation Path

Go to the OSPF Process Page—Redistribution Tab, page 67-38, then click the **Add** or **Edit** button beneath the Prefix Mapping table.

Related Topics

- OSPF Redistribution Mapping Dialog Box, page 67-39
- Redistributing Routes into OSPF, page 67-22

Field Reference

Table 67-19 OSPF Max Prefix Mapping Dialog Box

Element	Description
Process ID	The OSPF process into which other routes are being redistributed. The list contains the OSPF processes defined in the OSPF Process Page—Setup Tab, page 67-35.
Max Prefix	The maximum number of prefixes (routes) that can be redistributed into the selected OSPF process. Limiting the number of redistributed routes helps prevent the router from being flooded by an excessive number of routes.
Threshold	The percentage of the maximum prefix value that acts as a threshold for triggering warning messages. The default is 75%. **Note** This warning is triggered whether or not the Warning-Only check box is selected.
When maximum routes reached	The action to take when the maximum number of redistributed routes is reached: • Enforce Maximum Route—Prevents additional routes from being redistributed when the defined maximum prefix value is reached. This is the default. • Warning Only—Issues a warning when the maximum number of routes is reached, but does not prevent additional routes from being redistributed.

RIP Routing on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

Routing Information Protocol (RIP) is an Interior Gateway Protocol (IGP) that was created for use in small, homogeneous networks. RIP is a distance-vector protocol that sends routing-update messages at regular intervals (in a process called *advertising*) and whenever the network topology changes. When a router receives a routing update that includes changes to an entry, it updates its routing table to reflect the new route. If a router does not receive an update from another router for 180 seconds or more, it marks the routes served by the non-updating router as being unusable. If there is still no update after 240 seconds, the router removes all routing table entries for the non-updating router. Routing information is exchanged using UDP packets.

RIP evaluates routes by measuring the number of hops (the number of routers traversed) from the source to the destination. A directly connected network has a metric of zero. The maximum hop count allowed by RIP is 15. Any route with a hop count greater than 15 is considered unreachable.

Security Manager supports RIP version 2 only, which is described in RFC 1723. RIP 2 improves on the original RIP by enabling RIP messages to carry more information, which permits the use of a simple authentication mechanism (clear text or MD5) to secure table updates. RIP 2 also supports subnet masks, a critical feature that was not available in the original version of RIP.

The following topics describe the tasks you perform to create a RIP routing policy:

- Defining RIP Setup Parameters, page 67-42
- Defining RIP Interface Authentication Settings, page 67-43
- Redistributing Routes into RIP, page 67-44

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- OSPF Routing on Cisco IOS Routers, page 67-19
- EIGRP Routing on Cisco IOS Routers, page 67-8
- BGP Routing on Cisco IOS Routers, page 67-1

Defining RIP Setup Parameters

You configure RIP setup parameters by selecting the networks to include in the route and deciding whether any interfaces should be passive. These interfaces do not send routing updates to their neighbors. Additionally, you can enable auto-summarization, which reduces the size and complexity of the routing tables the router must maintain.

Related Topics

- Defining RIP Interface Authentication Settings, page 67-43
- Redistributing Routes into RIP, page 67-44
- RIP Routing on Cisco IOS Routers, page 67-42

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > RIP** from the Policy selector, then click the **Setup** tab in the work area.
- (Policy view) Select **Router Platform > Routing > RIP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Setup** tab.

The RIP Setup tab is displayed (see RIP Page—Setup Tab, page 67-45).

Step 2 Enter the addresses of the directly connected networks whose interfaces are to receive RIP updates. You can use a combination of addresses and network/host objects; separate addresses with commas. Click **Select** to select network/host objects from a list of existing objects, or to create new network/host objects. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 3 Enter the addresses of the passive interfaces, which are interfaces that should not send routing updates to their neighbors, if any. These interfaces continue to receive RIP routing broadcasts, which they use to populate their routing tables. Enter the names of one or more interfaces or interface roles; separate addresses with commas. Click **Select** to select interface names or roles from a list of existing objects, or to create new interface role objects. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 (Optional) Select the **Auto Summary** check box to enable the automatic summarization of subnet routes into network-level routes. Summarization reduces the size of routing tables, thereby reducing the complexity of the network.

Disable automatic summarization when you perform routing between disconnected subnets. When automatic summarization is turned off, subnets are advertised.

Defining RIP Interface Authentication Settings

You define neighbor authentication settings for RIP interfaces by selecting the interfaces and then selecting an authentication type, either MD5 or clear text.

Related Topics

- Defining RIP Setup Parameters, page 67-42
- Redistributing Routes into RIP, page 67-44
- RIP Routing on Cisco IOS Routers, page 67-42

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > RIP** from the Policy selector, then click the **Authentication** tab in the work area.
- (Policy view) Select **Router Platform > Routing > RIP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Authentication** tab.

The RIP Authentication tab is displayed. See Table 67-21 on page 67-46 for a description of the fields on this tab.

Step 2 On the RIP Authentication tab, select an interface definition from the table, then click **Edit**, or click **Add** to create a definition. The RIP Authentication dialog box appears. See Table 67-22 on page 67-47 for a description of the fields in this dialog box.

Step 3 Enter the name of the interface or interface role for which authentication is defined, or click **Select** to select an interface role from a list or to create a new one. For more information, see Specifying Interfaces During Policy Definition, page 6-76.

Step 4 Define interface authentication (MD5 or clear text).

Note We do not recommend that you use clear text authentication in RIP packets, because the unencrypted authentication key is sent in every packet. Use plain text authentication only when security is not an issue, for example, to ensure that misconfigured hosts do not participate in routing.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box. The defined interface appears on the RIP Authentication tab.

Redistributing Routes into RIP

Redistribution refers to using a routing protocol, such as RIP, to advertise routes that are learned by some other means, such as a different routing protocol, static routes, or directly connected routes. For example, you can redistribute routes from the OSPF routing protocol into your RIP route. Redistribution is necessary in networks that operate in multiple-protocol environments and can be applied to all IP-based routing protocols.

When you redistribute into RIP, you can maintain the original metric of the route by redistributing it transparently.

Before You Begin

- Define at least one RIP route. See Defining RIP Setup Parameters, page 67-42.

Related Topics

- Defining RIP Setup Parameters, page 67-42
- Defining RIP Interface Authentication Settings, page 67-43
- RIP Routing on Cisco IOS Routers, page 67-42

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > RIP** from the Policy selector, then click the **Redistribution** tab in the work area.
- (Policy view) Select **Router Platform > Routing > RIP** from the Policy Type selector. Select an existing policy or create a new one, and then click the **Redistribution** tab.

The RIP Redistribution tab is displayed. See Table 67-23 on page 67-48 for a description of the fields on this tab.

Step 2 On the RIP Redistribution tab, select a row from the RIP Redistribution Mappings table, then click **Edit**, or click **Add** to create a mapping. The RIP Redistribution Mapping dialog box appears. See Table 67-24 on page 67-49 for a description of the fields in this dialog box.

Step 3 Select the protocol whose routes you want to redistribute into RIP.

Note You can create a single mapping for each static route, BGP AS, EIGRP AS, and OSPF process.

Step 4 Define the metric (cost) of the redistributed routes by doing one of the following:

- Select the **Default Metric** check box, then enter the default metric of the redistributed routes. The metric determines the priority of the routes.
- Select the **Transparent** check box to maintain the original metric of the routes being redistributed into RIP.

Step 5 Click **OK** to save your definitions locally on the client and close the dialog box. The redistribution mapping appears in the Redistribution Mapping table on the RIP Redistribution tab.

RIP Routing Policy Page

RIP is a distance-vector routing protocol that uses hop count as the metric for path selection. Security Manager supports RIP version 2 only, which includes support for neighbor authentication when routing updates are exchanged.

You can configure RIP routing policies from the following tabs on the RIP Routing page:

- RIP Page—Setup Tab, page 67-45
- RIP Page—Authentication Tab, page 67-46
- RIP Page—Redistribution Tab, page 67-48

For more information, see RIP Routing on Cisco IOS Routers, page 67-42.

Navigation Path

- (Device view) Select **Platform > Routing > RIP** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > RIP** from the Policy Type selector. Right-click **RIP** to create a policy, or select an existing policy from the Shared Policy selector.

RIP Page—Setup Tab

Use the RIP Setup tab to create, edit, and delete RIP routes.

Navigation Path

Go to the RIP Routing Policy Page, page 67-45, then click the **Setup** tab.

Related Topics

- Defining RIP Setup Parameters, page 67-42
- RIP Page—Authentication Tab, page 67-46
- RIP Page—Redistribution Tab, page 67-48
- Specifying IP Addresses During Policy Definition, page 6-87
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 67-20 RIP Setup Tab

Element	Description
Networks	The directly connected networks associated with the RIP route. Enter one or more network addresses or network/host objects, separated by commas. Click **Select** to select network/host objects from a list of existing objects, or to create new objects.
Passive Interfaces	The interfaces that do not send updates to their routing neighbors. Enter one or more interface names or roles, separated by commas. Click **Select** to select interface names or roles from a list of existing objects, or to create new interface role objects.
Auto-Summary	When selected, enables the automatic summarization of subnet routes into network-level routes. Summarization reduces the size of routing tables, thereby reducing the complexity of the network. When deselected, automatic summarization is disabled. **Note** Disable automatic summarization when performing routing between disconnected subnets. When this feature is disabled, subnets are advertised.

RIP Page—Authentication Tab

Use the RIP Authentication tab to view, create, edit, and delete the neighbor authentication settings of RIP interfaces.

Navigation Path

Go to the RIP Routing Policy Page, page 67-45, then click the **Authentication** tab.

Related Topics

- Defining RIP Interface Authentication Settings, page 67-43
- RIP Page—Setup Tab, page 67-45
- RIP Page—Redistribution Tab, page 67-48
- RIP Routing Policy Page, page 67-45
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-21 RIP Authentication Tab

Element	Description
Interfaces	The name of an interface (as defined by an interface role) on which RIP is enabled.
Authentication	The type of RIP neighbor authentication that is enabled for the selected interface role—clear text or MD5.

Table 67-21 RIP Authentication Tab (continued)

Element	Description
Key ID	The identification number of the authentication key used for MD5 authentication.
Add button	Opens the RIP Authentication Dialog Box, page 67-47. From here you can define authentication for an additional RIP interface.
Edit button	Opens the RIP Authentication Dialog Box, page 67-47. From here you can edit the authentication properties of the selected RIP interface.
Delete button	Deletes the selected authentication definitions from the table.

RIP Authentication Dialog Box

Use the RIP Authentication dialog box to add or edit the neighbor authentication properties of RIP interfaces.

Navigation Path

Go to the RIP Page—Authentication Tab, page 67-46, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining RIP Interface Authentication Settings, page 67-43

Field Reference

Table 67-22 RIP Authentication Dialog Box

Element	Description
Interface	The interface for which you want to define authentication properties. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one. **Note** You cannot specify two different authentication configurations for the same interface.
Authentication	The type of authentication to apply to the interface: • MD5—(Recommended) Uses the MD5 hash algorithm for authentication. • Clear Text—Uses clear text for authentication. **Note** Use plain text authentication only when security is not an issue, for example, to ensure that misconfigured hosts do not participate in routing.
Key ID	Available only when MD5 is selected as the authentication type. The identification number of the authentication key. This number must be shared with all other devices sending updates to, and receiving updates from, the selected device. Valid values range from 0 to 2147483647.

Table 67-22 RIP Authentication Dialog Box (continued)

Element	Description
Key	The shared key used for authentication (MD5 or clear text). This key must be shared with all other devices sending updates to, and receiving updates from, the selected device. The key can contain up to 80 alphanumeric characters; the first character cannot be a number. Spaces are allowed. Enter the key again in the Confirm field.

RIP Page—Redistribution Tab

Use the RIP Redistribution tab to view, create, edit, and delete redistribution settings when performing redistribution into an RIP routing domain.

Note You must define RIP setup parameters before you can access the RIP Redistribution tab. See RIP Page—Setup Tab, page 67-45.

Navigation Path

Go to the RIP Routing Policy Page, page 67-45, then click the **Redistribution** tab.

Related Topics

- Redistributing Routes into RIP, page 67-44
- RIP Page—Authentication Tab, page 67-46
- Filtering Tables, page 1-48

Field Reference

Table 67-23 RIP Redistribution Tab

Element	Description
Protocol	The protocol that is being redistributed.
AS/Process ID	The autonomous system (AS) number or process ID of the route being redistributed.
Metric	The value that determines the priority of the redistributed route.
Match	When redistributing an OSPF process, indicates which types of OSPF routes are being redistributed.
Add button	Opens the RIP Redistribution Mapping Dialog Box, page 67-49. From here you can define a RIP redistribution mapping.
Edit button	Opens the RIP Redistribution Mapping Dialog Box, page 67-49. From here you can edit the selected RIP redistribution mapping.
Delete button	Deletes the selected redistribution mappings from the table.

RIP Redistribution Mapping Dialog Box

Use the RIP Redistribution Mapping dialog box to add or edit the properties of an RIP redistribution mapping.

Navigation Path

Go to the RIP Page—Redistribution Tab, page 67-48, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Redistributing Routes into RIP, page 67-44

Field Reference

Table 67-24 RIP Redistribution Mapping Dialog Box

Element	Description
Protocol to Redistribute	The routing protocol that is being redistributed: • Static—Redistributes static routes. You can define a single mapping for each route. • EIGRP—Redistributes an EIGRP autonomous system. Enter the AS number in the displayed field. You can define a single mapping for each AS. • BGP—Redistributes a BGP autonomous system. You can define a single BGP mapping on each device. If you configured a BGP AS in the BGP Setup tab, the AS number is displayed. Otherwise, a message is displayed indicating that no BGP AS was defined. See BGP Page—Redistribution Tab, page 67-7.
Protocol to Redistribute (continued)	• OSPF—Redistributes a different OSPF process. You can define a single mapping for each process. Select a process from the displayed list, then select one or more match criteria: – Internal—Routes that are internal to a specific AS. – External1—Routes that are external to the AS and imported into OSPF as a Type 1 external route. – External2—Routes that are external to the AS and imported into the selected process as a Type 2 external route. – NSAAExternal1—Not-So-Stubby Area (NSSA) routes that are external to the AS and imported into the selected process as Type 1 external routes. – NSAAExternal2—(NSSA) routes that are external to the AS and imported into the selected process as Type 2 external routes. • Connected—Redistributes routes that are established automatically by virtue of having enabled IP on an interface. These routes are redistributed as external to the AS.
Default Metric	Establishes a default value for the redistributed route. Valid values range from 0 to 16.

Table 67-24 RIP Redistribution Mapping Dialog Box (continued)

Element	Description
Transparent Metric	When selected, maintains the original metric of the route being redistributed. When deselected, the value specified in the Metric field is used.

Static Routing on Cisco IOS Routers

Note From version 4.17, though Cisco Security Manager continues to support IOS features/functionality, it does not support any bug fixes or enhancements.

You can configure static routing policies to ensure that the router correctly forwards packets to their destination when a route cannot be built dynamically. By default, static routes have a default administrative distance of 1 (implying a directly connected network), which causes them to override any dynamic routes discovered for the same host or network. You can, however, define a larger administrative distance to a static route so that it does not take precedence over a corresponding dynamic route.

For example, EIGRP routes have a default administrative distance of 5. To have a static route that can be overridden by an EIGRP route, you must specify an administrative distance greater than 5. This feature is useful when you define the static route as a "floating" route, which is inserted into the routing table only when the preferred route is unavailable.

Tip When you use the static route as a backup, "floating" route, specify the interface through which the next hop IP address can be reached instead of entering a specific IP address. Otherwise, the "floating" route is not inserted in the routing table if the primary link fails. For more information, see *Specifying a Next Hop IP Address for Static Routes* on Cisco.com at this URL: http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00800ef7b2.shtml

Related Topics

- Defining Static Routes, page 67-50

Defining Static Routes

To define a static route, you must define the IP address (and optionally, the metric) of the hop gateway to which the router forwards packets destined to the selected host or network. You can define as many static routes as required.

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- RIP Routing on Cisco IOS Routers, page 67-42
- OSPF Routing on Cisco IOS Routers, page 67-19
- EIGRP Routing on Cisco IOS Routers, page 67-8
- BGP Routing on Cisco IOS Routers, page 67-1

Step 1 Do one of the following:

- (Device view) Select **Platform > Routing > Static Routing** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > Static Routing** from the Policy Type selector. Select an existing policy or create a new one.

The Static Routing page is displayed. See Table 67-25 on page 67-52 for a description of the fields on this page.

Step 2 On the Static Routing page, select a static route from the table, then click **Edit**, or click **Add** to create a route. The Static Routing dialog box appears. See Table 67-26 on page 67-53 for a description of the fields in this dialog box.

Step 3 (Optional) Select the **Use as Default Route** check box to make this route the default route for all unknown outbound packets.

Step 4 In the Prefix field, enter the address for the destination network, or click **Select** to select a network/host object from a list or to create a new one. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 5 Select a forwarding option:

- To define the router interface that forwards packets to the remote network, select **Forwarding Interface** and enter the name of an interface or interface role. You can click Select to select an interface role from a list or to create a new one. See Understanding Interface Role Objects, page 6-73 and Selecting Objects for Policies, page 6-2.
- To specify the next hop router that receives and forwards packets to the remote network, select **Forwarding IP**, then enter the address in the field provided, or click **Select** to select a network/host object from a list or to create a new one. For more information, see Specifying IP Addresses During Policy Definition, page 6-87.

Step 6 (Optional) In the Distance Metric field, enter the number of hops to the next hop address for this router. This metric identifies the priority of the static route. If two routing entries specify the same network, the route with the lower metric value (that is, the lower cost) is given a higher priority and is selected.

If no value is specified, the default is 1, which implies a directly connected network.

Step 7 (Optional) Select the **Permanent route** check box to prevent this static route entry from being deleted, even in cases in which the interface is shut down or the router cannot communicate with the next router.

Step 8 Click **OK** to save your definitions locally on the client and close the dialog box. The static route appears in the table on the Static Routing page.

Static Routing Policy Page

Use the Static Routing page to create, edit, and delete static routes. For more information, see Defining Static Routes, page 67-50.

Navigation Path

- (Device view) Select **Platform > Routing > Static Routing** from the Policy selector.
- (Policy view) Select **Router Platform > Routing > Static Routing** from the Policy Type selector. Right-click **Static Routing** to create a policy, or select an existing policy from the Shared Policy selector.

Related Topics

- Static Routing on Cisco IOS Routers, page 67-50
- Table Columns and Column Heading Features, page 1-49
- Filtering Tables, page 1-48

Field Reference

Table 67-25 Static Routing Page

Element	Description
Prefix	The destination IP address of the static route.
Prefix Mask	The net mask of the selected IP address.
Default Route	Indicates whether the static route is the default route for unknown packets being forwarded by this router.
Interface or IP Address	The IP address or the interface name associated with the gateway router that is the next hop address for this router.
Distance	The number of hops from the gateway IP to the destination. The metric determines the priority of this route. The fewer the hops, the higher the priority assigned to the route, based on lower costs. When two routing entries specify the same network, the entry with the lower metric (that is, the higher priority) is selected.
Permanent Route	Indicates whether the static route is defined as a permanent route, which means that it will not be removed even if the interface is shut down or if the router is unable to communicate with the next router.
Add button	Opens the Static Routing Dialog Box, page 67-52. From here you can create a static route.
Edit button	Opens the Static Routing Dialog Box, page 67-52. From here you can edit the selected static route.
Delete button	Deletes the selected static routes from the table.

Static Routing Dialog Box

Use the Static Routing dialog box to add or edit static routes.

Navigation Path

Go to the Static Routing Policy Page, page 67-51, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Defining Static Routes, page 67-50
- Static Routing on Cisco IOS Routers, page 67-50

Field Reference

Table 67-26 Static Routing Dialog Box

Element	Description
Destination Network	Address information for the destination network defined by this static route. • Use as Default Route—When selected, makes this the default route on this router. A default route is used when the route from a source to a destination is unknown or when it is not feasible for the router to maintain many routes in its routing table. All unknown outbound packets are forwarded over the default route. When deselected, this static route is not the default route. • Prefix—The IP address of the destination network. Enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one. The prefix must be a class A, B, or C network or host IP. A host IP can begin with 0 unless it contains a discontiguous mask. All subnet addresses are valid.
Forwarding (Next Hop)	The method of forwarding data to the destination network: • Forwarding Interface—The router interface that forwards packets to the remote network. Enter the name of an interface or interface role, or click **Select** to select the object from a list or to create a new one. • Forwarding IP—The IP address of the next hop router that receives and forwards packets to the remote network. Enter an IP address or the name of a network/host object, or click **Select** to select the object from a list or to create a new one.
Distance Metric	The number of hops to the destination network (gateway IP). The default is 1 if no value is specified. The range is from 1 to 255. This metric (also known as *administrative distance*) is a measurement of route expense based on the number of hops to the network on which a specified host resides. This hop count includes all the networks a packet must traverse, including the destination network. Therefore, all directly connected networks have a metric of 1. Because the metric is based on expense, it is used to identify the priority of the static route. If two routing entries specify the same network, the route with the lower metric value (that is, the lower cost) is given a higher priority and is selected. **Note** Under certain circumstances, it is useful to assign a static route a lower priority (larger distance metric) than a dynamic route. This enables the static route to act as a backup, "floating," route when the dynamic route is unavailable.
Permanent route	When selected, prevents this static route entry from being deleted, even in cases where the interface is shut down or the router cannot communicate with the next router. When deselected, this static route can be deleted.

CHAPTER 68

Managing Cisco Catalyst Switches and Cisco 7600 Series Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any bug fixes or enhancements.

Cisco Security Manager supports the management and configuration of security services and other platform-specific services on Cisco Catalyst switches and Cisco 7600 Series routers.

You can manage Catalyst switches and 7600 devices configured in VTP transparent or VTP client/server mode. Security Manager manages switches configured in client/server mode by bypassing VLAN database management on the device (including VLAN creation, deletion, and monitoring VLANs in the VLAN database on switches).

This chapter contains the following topics:

- Discovering Policies on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-1
- Viewing Catalyst Summary Information, page 68-2
- Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups, page 68-3
- Interfaces, page 68-5
- VLANs, page 68-25
- VLAN Groups, page 68-31
- VLAN ACLs (VACLs), page 68-36
- IDSM Settings, page 68-43

Discovering Policies on Cisco Catalyst Switches and Cisco 7600 Series Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any bug fixes or enhancements.

You can discover the configurations of your Cisco Catalyst switches and Cisco 7600 Series Routers (as well as the configurations of the services modules and security contexts associated with them) and import the configurations as policies into Security Manager. This makes it possible to add existing devices and manage them with Security Manager without having to configure each device manually, policy by policy. For more information, see Adding Devices to the Device Inventory, page 3-7.

You can discover any command that Security Manager can configure. Discovery ignores unsupported commands, which means that they are left intact on the device even after subsequent deployments. Additionally, in cases where Security Manager can discover the command, but not all the subcommands and keywords related to that command, the unsupported elements are ignored and left intact on the device.

At any time, you can also *rediscover* the configurations of devices that you are already managing with Security Manager. Be aware, however, that we do not recommend rediscovery generally because performing rediscovery overwrites the policies that you have defined in Security Manager. For more information, see Discovering Policies on Devices Already in Security Manager, page 5-15.

Note We recommend that you perform deployment immediately after you discover policies, *before* you make any changes to policies or unassign policies from the device. (This recommendation also applies to any services module or security context hosted by the device.) Otherwise, the changes that you configure in Security Manager might not be deployed to the device. See Working with Deployment and the Configuration Archive, page 8-25.

Related Topics

- Understanding Policies, page 5-1
- Discovering Policies, page 5-12
- Chapter 68, “Managing Cisco Catalyst Switches and Cisco 7600 Series Routers”
- Working with Deployment and the Configuration Archive, page 8-25

Viewing Catalyst Summary Information

Use the Catalyst Summary Info page to view high-level system information, including any service modules, ports, and VLANs that Security Manager has discovered.

To view Catalyst summary information, in Device view, right-click a Catalyst switch or Cisco 7600 Series router, then select **Catalyst Summary Info**, or select **Tools > Catalyst Summary Info**.

Note If Security Manager has not completed discovery for a particular Cisco Catalyst switch or Cisco 7600 Series router, the Catalyst Summary Info page for that device displays this message: “No information is available. This information is acquired during device discovery.”

Related Topics

- IDSM Settings Page, page 68-47
- VLAN Access Lists Page, page 68-39
- Filtering Tables, page 1-48

Field Reference

Table 68-1 Catalyst Summary Info Page

Element	Description
Hostname	Displays the configured hostname of the device.
Device Type	Displays the device type.
Serial Number	Displays the serial number of the device.
OS Version	Displays the Cisco IOS image version the device is running.
Image	Displays the name of the image running on the device.
Last Update	Displays a time stamp for the most recent discovery.
Total Ports	Displays the total number of configured ports, combining access ports, routed ports, and trunk ports.
Access Ports	Displays the number of configured access ports on the chassis.
Trunk Ports	Displays the number of configured trunk ports on the chassis.
Routed Ports	Displays the number of configured routed ports on the chassis.
Total VLANs	Displays the total number of configured VLANs on the chassis and all its services modules.
Layer 2 VLANs	Displays the number of VLANs that run on Layer 2.
Layer 3 VLANs	Displays the number of VLANs that run on Layer 3.
Service Module Table	
Slot	Identifies the slot to which a service module is attached.
Device Type	Displays a brief description of the service module.
Serial Number	Displays the serial number of the service module.
Model	Displays the model type of the service module.
OS Version	Identifies the OS version that is installed and running on the service module.
Assigned VLANs	Displays the total number of VLANs to which an FWSM is assigned. **Tip** Click the Summary tab of the Interfaces/VLANs policy to learn which VLANs are assigned to an IDSM or a VPNSM.
Contexts	Displays the total number of configured security contexts for an FWSM that runs in multicontext mode. **Tip** Click the Summary tab of the Interfaces/VLANs policy to learn how many virtual sensors are configured for an IDSM.

Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups

Use the Summary tab of the Interfaces/VLAN policy to view attributes of all VLANs, VLAN groups, interfaces, and subinterfaces configured on supported Catalyst 6500 Series and 7600 Series chassis and their associated services modules.

To view summary interface information, in Device view, select **Interfaces/VLANs** from the Policy selector, then click the **Summary** tab.

Note The Summary tab is available only for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers.

Related Topics

- Interfaces/VLANs Page—VLANs Tab, page 68-27
- Interfaces/VLANs Page—VLAN Groups Tab, page 68-33
- Interfaces/VLANs Page—Interfaces Tab, page 68-8
- Viewing Catalyst Summary Information, page 68-2
- Filtering Tables, page 1-48

Field Reference

Table 68-2 Interfaces/VLANs Page—Summary Tab

Element	Description
VLAN ID	The VLAN ID associated with an interface or subinterface. The VLAN ID specifies where 802.1Q tagged packets are sent and received on the specified an interface or subinterface; without a VLAN ID, the interface or subinterface cannot send or receive traffic. **Note** All VLAN IDs must be unique among all subinterfaces configured on the same physical interface.
VLAN Name	Name of the VLAN that corresponds to an interface or subinterface. For example, VLAN003 or Trunk1.
VLAN Group	Numeric identity of a VLAN group that is configured on the VLAN that a table row describes.
VLAN Type	Specifies whether a VLAN has access to Layer 2 or Layer 3.
IP Address/Mask	The IP address and corresponding subnet mask of the VLAN configured on an interface or subinterface.
Access Port	Displays the assigned name, if a name is assigned, of the access port that a VLAN uses.
Trunk Port	Specifies which VLANs are permitted to carry traffic over the trunk.
Slot (-Port)	Associates the chassis slot number (in which the relevant services module is installed) with the port number, as a hyphenated pair in the format *x* -*y*, for example 3-1.
Blade Type	Identifies the kind of services module on which a particular VLAN is configured, such as FWSM or VPNSM.
Security Context	Identifies the security context associated with an interface, but only if Multiple Mode is active on the installed module and an Admin context is configured for the module.

Table 68-2 Interfaces/VLANs Page—Summary Tab (continued)

Element	Description
Security Context Interface	Displays the physical interface and subinterface IDs for which a security context inspects traffic. The displayed ID can represent a physical interface, a single sub-interface (defined as a range of one), or a range of sub-interfaces.
Security Level	Displays the security level of an interface, where values range from 0 (the lowest security) to 100 (the highest): • For an outside interface, the default is 0. • For an inside interface, the default is 100. • For an interface in the DMZ, the default is typically from 1 to 99.

Interfaces

You use the Interfaces tab on the Interfaces/VLANs page to view and manage the following types of ports:

- Access ports—A switching port that is used to connect host machines or servers. An access port belongs to and carries the traffic of only one VLAN. Traffic is received and sent in native formats with no VLAN tagging.
- Trunk ports—A switching port operating at Layer 2 to carry the traffic of multiple VLANs. Traffic is tagged with a VLAN number to differentiate traffic from each VLAN. A trunk port is used to connect switches to switches or to connect switches to routers.
- Routed ports—A physical port that acts like a port on a router. A routed port is not associated with a particular VLAN, and it behaves like a regular router interface. You can configure a routed port with a Layer 3 routing protocol.
- Dynamic ports—A port that can change dynamically to a trunk port if the neighboring port is configured as a trunk port.
- Unsupported ports—Ports on the Catalyst device that are not supported by Security Manager.

To display the Interfaces tab, select a Catalyst device in Device view, select **Interfaces/VLANs** from the Policy selector, then click the **Interfaces** tab in the work area.

The following topics describe the actions you can perform when defining interfaces on Catalyst devices:

- Creating or Editing Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-6
- Deleting Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-7
- Interfaces/VLANs Page—Interfaces Tab, page 68-8

Related Topics

- VLANs, page 68-25
- VLAN Groups, page 68-31
- VLAN ACLs (VACLs), page 68-36
- Chapter 68, “Managing Cisco Catalyst Switches and Cisco 7600 Series Routers”

Creating or Editing Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any bug fixes or enhancements.

You can create access ports, routed ports, or trunk ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, with these restrictions:

- Each interface must have a name.
- You can associate an access port with only one VLAN.
- You can associate a trunk port with one or more VLANs.

Related Topics

- Deleting Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-7
- Creating or Editing VLANs, page 68-26
- Creating or Editing VLAN Groups, page 68-32
- Interfaces/VLANs Page—Interfaces Tab, page 68-8
- Interfaces, page 68-5

Step 1 (Device view) Select a Catalyst device, select **Interfaces/VLANs** from the Policy selector, then click the Interfaces tab in the work area.

The Interfaces tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—Interfaces Tab, page 68-8.

Step 2 Do one of the following:

- To define the attributes of a new interface, click **Add Row**.
- To edit the attributes of an interface, select it in the list, then click **Edit Row**.

Step 3 (Optional) Deselect the **Enable Interface** check box if you want this interface to be in shutdown mode.

Step 4 From the Type list, select **Interface** or **Subinterface**:

Step 5 (Interfaces only) Enter a name for the interface. You can click **Select** to open a dialog box that will help you generate a standard name based on interface type and details about the interface's location, such as card, slot, and subinterface. For more information on using the dialog box to generate an interface name, see Interface Auto Name Generator Dialog Box, page 62-12.

Step 6 (Interfaces only) Select an option from the **Mode** list to specify the port configuration type. The fields in the dialog box vary according to your selection.

Step 7 (Subinterfaces only) Select the parent interface of the subinterface, then enter the ID number.

Step 8 Define or configure the settings for the type that you selected:

- Access Port—See Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9 for a description of the fields.
- Routed Port—See Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12 for a description of the fields.

- Trunk Port—See Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14 for a description of the fields.
- Dynamic Port—See Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18 for a description of the fields.
- Subinterface—See Create and Edit Interface Dialog Boxes—Subinterfaces, page 68-22 for a description of the fields.
- Unsupported—See Create and Edit Interface Dialog Boxes—Unsupported Mode, page 68-24 for a description of the fields.

Step 9 From the **Speed** list, select an option to define the speed of the interface.

Step 10 If you defined a specific speed for the interface, and therefore the Duplex list is enabled, select a duplexing option.

Step 11 In the MTU field, enter the maximum transmission unit value.

Step 12 Configure whether to use flow control on inbound (Receive) and outbound (Send) traffic.

Step 13 (Optional) Enter a description for the interface in the **Description** field.

Step 14 Click **OK** to save your definitions locally on the client and close the dialog box.

Deleting Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers

Note From version 4.17, though Cisco Security Manager continues to support Cisco Catalyst switches features/functionality, it does not support any bug fixes or enhancements.

Although you can delete the definition of an interface at any time, use this option with great care. If the relevant device includes the interface definition in any policy definitions, deleting the interface causes these policy definitions to fail when they are deployed to the device.

Related Topics

- Creating or Editing Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-6
- Interfaces, page 68-5

Step 1 (Device view) Select a Cisco Catalyst switch or Cisco 7600 Series router from the Device selector.

Step 2 Select **Interfaces/VLANs** from the Policy selector.

Step 3 Click the Interfaces tab in the work area.

The Interfaces tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—Interfaces Tab, page 68-8.

Step 4 Select an interface from the table, then click **Delete Row**. The interface is deleted.

Interfaces/VLANs Page—Interfaces Tab

Use the Interfaces tab to view and configure interfaces and subinterfaces on supported Cisco Catalyst switches and Cisco 7600 Series routers and their associated services modules (blades).

Navigation Path

(Device view) Select **Interfaces/VLANs** from the Device selector, then click the **Interfaces** tab.

Related Topics

- Interfaces/VLANs Page—VLANs Tab, page 68-27
- Interfaces/VLANs Page—VLAN Groups Tab, page 68-33
- Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups, page 68-3
- Filtering Tables, page 1-48

Field Reference

Table 68-3 Interfaces/VLANs Page—Interfaces Tab

Element	Description
Name	Interface type, chassis slot, and the number of the interface card. For example, *FastEthernet 2/7* means Fast Ethernet, slot 2, interface 7.
Mode	Configuration mode for physical ports: • Access • Routed • Trunk • Dynamic Auto • Dynamic Desirable • Unsupported
VLAN ID	The VLAN ID associated with the described subinterface, displayed only for Ethernet interfaces and VLAN interfaces.
IP Address	The IP address of the interface.
Enabled	Indicates whether the interface is enabled or disabled (shutdown state).
Interface Roles	The interface roles whose naming patterns match this interface. See Understanding Interface Role Objects, page 6-73.
Description	An optional description of the interface.

Table 68-3 Interfaces/VLANs Page—Interfaces Tab (continued)

Element	Description
Add Row button	Opens the Create Interface dialog box, where you can define a new interface. For more information, see the instructions for the relevant mode: • Access Port Mode—Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9. • Routed Port Mode—Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12 • Trunk Port Mode—Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14 • Dynamic Mode—Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18
Edit Row button	Opens the Edit Interface dialog box, where you can edit the selected interface. For more information, see the instructions for the relevant mode: • Access Port Mode—Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9. • Routed Port Mode—Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12 • Trunk Port Mode—Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14 • Dynamic Mode—Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18 • Unsupported—Create and Edit Interface Dialog Boxes—Unsupported Mode, page 68-24
Delete Row button	Deletes the selected interface.

Create and Edit Interface Dialog Boxes—Access Port Mode

Use the Create Interface dialog box (or the Edit Interface dialog box) to configure the attributes of physical and virtual interfaces that run in access port mode.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, click **Add** or **Edit** to open the Create/Edit Interface dialog box, then select **Access Port** from the Mode list.

Related Topics

- Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12
- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14
- Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18
- Interface Auto Name Generator Dialog Box, page 62-12
- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 68-4 Create and Edit Interface Dialog Boxes—Access Port Mode

Element	Description
Enable Interface	When selected, enables the interface. When deselected, disables the interface using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface. For details about defining a subinterface, see Create and Edit Interface Dialog Boxes—Subinterfaces, page 68-22.
Name (Select button)	Displays the generated interface name, if the name has been set. Click **Select** to open the Interface Auto Name Generator Dialog Box, page 62-12. From here, you can enter or edit the details that Security Manager uses to generate an interface name.
Mode	The port configuration type for this interface. Select **Access Port** to display the configuration options that are relevant for access ports.
Access Port settings	
VLAN ID (Select button)	Displays the interface-specific identity of the VLAN to use in access port mode, if you have selected a VLAN. Otherwise, click **Select** to open the VLAN Selector Dialog Box, page 68-35. The VLAN ID specifies where 802.1Q tagged packets are sent and received on the subinterface; without a VLAN ID, the subinterface cannot send or receive traffic. Valid values range from 1 to 4094. Some VLAN IDs might be reserved on connected devices, so see the device documentation for more information. For multiple context mode, you can only set the VLAN in the system configuration. **Note** All VLAN IDs must be unique among all subinterfaces configured on the same physical interface. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the device.
Enable Port Security	When selected, enables you to restrict input to an interface by limiting the MAC addresses that are allowed to access the port. When deselected, disables port security.
Max. MAC Addresses	Applies only when Enable Port Security is selected. The maximum number of secure MAC addresses for the interface. Valid values range from 1 to 4097. **Note** Secure MAC addresses are configured dynamically using the MAC addresses of connected devices.

Table 68-4 Create and Edit Interface Dialog Boxes—Access Port Mode (continued)

Element	Description
Violation Policy	The action to take if a security violation occurs: • Port Security Protect—Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses and the count drops below the maximum value. • Port Security Restrict—Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses and the count drops below the maximum value. In addition, it causes the SecurityViolation counter to increment. • Port Security Shutdown—Immediately puts the interface into the error-disabled state and sends an SNMP trap notification. A security violation occurs if a workstation whose MAC address is not in the address table attempts to access the interface after the maximum number of secure MAC addresses is configured.
Enable VACL Capture	When selected, enables VACL capture. If the capture bit is set, ports with the capture function enabled can receive forwarded packets. When deselected, disables VACL capture.
Capture VLANs (Select button)	Enables you to identify the VLANs where VACLs should receive forwarded VLAN packets. This option is available if you selected the Enable VACL Capture check box. Enter a comma-separated list of VLAN IDs or click **Select** to open the VLAN Selector Dialog Box, page 68-35. VACLs can capture VLAN packets only when they are initially routed or bridged into the VLAN. Only forwarded packets can be captured.
Common interface settings	
Speed	The speed of the physical interface: • 10—Transmits at 10 Mbps. • 100—Transmits at 100 Mbps. • 1000—Transmits at 1,000 Mbps. • 10000—Transmits at 10,000 Mbps. • Auto—If Speed is set to Auto, both Speed and Duplex are autonegotiated. • Non-Negotiate—Disables link negotiation.
Duplex	The duplex setting of the interface: • Auto—Autonegotiates the duplex. • Half—Sends and receives data, but not at the same time • Full—Sends and receives data at the same time. If the speed is set to Auto, the duplex setting must also be set to Auto.
MTU	The maximum transmission unit, which refers to the largest packet size (in bytes) that can be handled by the interface. The range of valid values depends on the interface type.

Table 68-4 Create and Edit Interface Dialog Boxes—Access Port Mode (continued)

Element	Description
Description	A text description of the interface. Enter up to 240 characters on a single line, without using carriage returns. **Note** For multiple context mode, the system description is independent of the context description.
Flow Control Receive	The flow control setting for incoming frames: • Off—The port does not use flow control, regardless of whether the neighboring port requests flow control. • On—The port uses flow control, as dictated by the neighboring port. • Desired—The port allows, but does not require, flow control frames. Flow control frames (also called pause frames) are special packets that signal a source to stop sending frames for a defined interval when buffers are full.
Flow Control Send	The flow control setting for outgoing frames: • Off—The port does not send flow control frames to the neighboring port. • On—The port sends flow control frames to the neighboring port. • Desired—The port allows, but does not require, flow control frames.
Roles	Lists the interface roles associated with the interface. Interface roles are objects that are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. See Understanding Interface Role Objects, page 6-73.

Create and Edit Interface Dialog Boxes—Routed Port Mode

Use the Create Interface dialog box (or the Edit Interface dialog box) to configure the attributes of physical interfaces that run in routed port mode on Layer 3.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, click **Add** or **Edit** to open the Create/Edit Interface dialog box, then select **Routed Port** from the Mode list.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14
- Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18
- Understanding Interface Role Objects, page 6-73
- Selecting Objects for Policies, page 6-2
- Understanding Networks/Hosts Objects, page 6-80

Field Reference

Table 68-5 Create and Edit Interface Dialog Boxes—Routed Port Mode

Element	Description
Enable Interface	When selected, enables the interface. When deselected, disables the interface using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface. For details about defining a subinterface, see Create and Edit Interface Dialog Boxes—Subinterfaces, page 68-22.
Name (Select button)	Displays the generated interface name, if the name has been set. Click **Select** to open the Interface Auto Name Generator Dialog Box, page 62-12. From here, you can enter or edit the details that Security Manager uses to generate an interface name.
Mode	The port configuration type for this interface. Select **Routed Port** to display the configuration options that are relevant for routed ports.
Routed Port settings	
IP Type	The type of IP address used by the port: • Static IP—Specifies that the interface uses a permanent IP address and activates related GUI elements.
IP Address (Select button)	Enables you to enter an IP address, or you can click **Select** to open the Networks/Hosts Selector, where you can select an IP address.
Helper IP Addresses (Select button)	Enables you to assign a helper IP address to the interface. A helper IP address converts broadcast DHCP requests to unicast requests that are directed exclusively to the DHCP server.
Mask	Enables you to specify the subnet mask. You can enter a netmask value or you can select a netmask from the list. If you enter a netmask, you can express its value in dotted decimal format (for example, 255.255.255.0) or you can enter the number of bits (for example, 24). **Note** Do not use 255.255.255.254 or 255.255.255.255 for any interface that is connected to your network; these netmasks cause all traffic on an interface to stop.
Common interface settings	
Speed	The speed of the physical interface: • 10—Transmits at 10 Mbps. • 100—Transmits at 100 Mbps. • 1000—Transmits at 1,000 Mbps. • 10000—Transmits at 10,000 Mbps. • Auto—If Speed is set to Auto, both Speed and Duplex are autonegotiated. • Non-Negotiate—Disables link negotiation.

Table 68-5 Create and Edit Interface Dialog Boxes—Routed Port Mode (continued)

Element	Description
Duplex	The duplex setting of the interface: • Auto—Autonegotiates the duplex. • Half—Sends and receives data, but not at the same time • Full—Sends and receives data at the same time. If the speed is set to Auto, the duplex setting must also be set to Auto.
MTU	The maximum transmission unit, which refers to the largest packet size (in bytes) that can be handled by the interface. The range of valid values depends on the interface type.
Description	A text description of the interface. Enter up to 240 characters on a single line, without using carriage returns. **Note** For multiple context mode, the system description is independent of the context description.
Flow Control Receive	The flow control setting for incoming frames: • Off—The port does not use flow control, regardless of whether the neighboring port requests flow control. • On—The port uses flow control, as dictated by the neighboring port. • Desired—The port allows, but does not require, flow control frames. Flow control frames (also called pause frames) are special packets that signal a source to stop sending frames for a defined interval when buffers are full.
Flow Control Send	The flow control setting for outgoing frames: • Off—The port does not send flow control frames to the neighboring port. • On—The port sends flow control frames to the neighboring port. • Desired—The port allows, but does not require, flow control frames.
Roles	Lists the interface roles associated with the interface. Interface roles are objects that are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. See Understanding Interface Role Objects, page 6-73.

Create and Edit Interface Dialog Boxes—Trunk Port Mode

Use the Create Interface dialog box (or the Edit Interface dialog box) to configure the attributes of physical and virtual interfaces that run in trunk port mode.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, click **Add** or **Edit** to open the Create/Edit Interface dialog box, then select **Trunk Port** from the Mode list.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12
- Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18
- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 68-6 Create and Edit Interface Dialog Boxes—Trunk Port Mode

Element	Description
Enable Interface	When selected, enables the interface. When deselected, disables the interface using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface. For details about defining a subinterface, see Create and Edit Interface Dialog Boxes—Subinterfaces, page 68-22.
Name (Select button)	Displays the generated interface name, if the name has been set. Click **Select** to open the Interface Auto Name Generator Dialog Box, page 62-12. From here, you can enter or edit the details that Security Manager uses to generate an interface name.
Mode	The port configuration type for this interface. Select **Trunk Port** to display the configuration options that are relevant for trunk ports.
Trunk Port settings	
Encapsulation	Select one of the following: • DOT1Q—Specifies VLAN encapsulation on the trunk link, as defined by the IEEE 802.1Q standard. Applies only to Ethernet subinterfaces. • ISL—Specifies ISL encapsulation on the trunk link. 10-Gigabit Ethernet ports do not support ISL encapsulation. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the router.

Table 68-6 Create and Edit Interface Dialog Boxes—Trunk Port Mode (continued)

Element	Description
Native VLAN (Select button)	Enables you to select the Native VLAN to associate with this interface, using the ID specified in the VLAN ID field. (If no VLAN ID is specified for the Native VLAN, the default is 1.) This option applies to you only if you are configuring a physical interface that is meant to serve as an 802.1Q trunk interface. You must first specify DOT1Q as the encapsulation type. The Native VLAN of a trunk interface is the VLAN to which all untagged VLAN packets are logically assigned. This includes the management traffic associated with the VLAN. When deselected, the Native VLAN is not associated with this interface. **Note** The Native VLAN cannot be configured on a subinterface of the trunk interface. Be sure to configure the same Native VLAN value at both ends of the link; otherwise, traffic may be lost or sent to the wrong VLAN. Click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can associate a native VLAN with the described interface.
Enable DTP negotiation	When selected, enables Dynamic Trunking Protocol (DTP) negotiation. DTP manages trunk auto-negotiation (ISL and 802.1Q) between devices. When deselected, disables DTP negotiation.
Allowed VLANs (Select button)	Enables you to specify which VLANs are allowed on the trunk. Enter the VLAN IDs. Use commas to separate multiple VLANs or use a hyphen to indicate a range of VLANs (for example, 12,17,22 or 2-200). Valid IDs range from 1 to 4094. Or, click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select the VLANs to include on the trunk.
Prune VLANs (Select button)	Enables you to specify which VLANs are eligible for pruning. Enter the VLAN IDs. Use commas to separate multiple VLANs or use a hyphen to indicate a range of VLANs (for example, 12,17,22 or 2-200.) Or, click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select the VLANs that are eligible for pruning.
Enable VACL Capture	When selected, enables VACL capture. If the capture bit is set, ports with the capture function enabled can receive forwarded packets. When deselected, disables VACL capture.
Capture VLANs (Select button)	Enables you to identify the VLANs where VACLs should receive forwarded VLAN packets. This option is available if you selected the Enable VACL Capture check box. Enter a comma-separated list of VLAN IDs, or click **Select** to open the VLAN Selector Dialog Box, page 68-35. VACLs can capture VLAN packets only when they are initially routed or bridged into the VLAN. Only forwarded packets can be captured.

Table 68-6 Create and Edit Interface Dialog Boxes—Trunk Port Mode (continued)

Element	Description
Enable Port Security	Applies only to devices running IOS Software Version 12.2(18)SXE2 or later. When selected, enables you to restrict input to an interface by limiting the MAC addresses that are allowed to access the port. When deselected, disables port security. **Note** If you select this option, the Enable DTP Negotiation option is automatically deselected.
Max. MAC Addresses	Applies only when Enable Port Security is selected. The maximum number of secure MAC addresses for the interface. Valid values range from 1 to 4097. **Note** Secure MAC addresses are configured dynamically using the MAC addresses of connected devices.
Violation Policy	The action to take if a security violation occurs: • Port Security Protect—Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses and the count drops below the maximum value. • Port Security Restrict—Drops packets with unknown source addresses until you remove a sufficient number of secure MAC addresses and the count drops below the maximum value. In addition, it causes the SecurityViolation counter to increment. • Port Security Shutdown—Immediately puts the interface into the error-disabled state and sends an SNMP trap notification. A security violation occurs if a workstation whose MAC address is not in the address table attempts to access the interface after the maximum number of secure MAC addresses is configured.
Common interface settings	
Speed	The speed of the physical interface: • 10—Transmits at 10 Mbps. • 100—Transmits at 100 Mbps. • 1000—Transmits at 1,000 Mbps. • 10000—Transmits at 10,000 Mbps. • Auto—If Speed is set to Auto, both Speed and Duplex are autonegotiated. • Non-Negotiate—Disables link negotiation.
Duplex	The duplex setting of the interface: • Auto—Autonegotiates the duplex. • Half—Sends and receives data, but not at the same time • Full—Sends and receives data at the same time. If the speed is set to Auto, the duplex setting must also be set to Auto.

Table 68-6 Create and Edit Interface Dialog Boxes—Trunk Port Mode (continued)

Element	Description
MTU	The maximum transmission unit, which refers to the largest packet size (in bytes) that can be handled by the interface. The range of valid values depends on the interface type.
Description	A text description of the interface. Enter up to 240 characters on a single line, without using carriage returns. **Note** For multiple context mode, the system description is independent of the context description.
Flow Control Receive	The flow control setting for incoming frames: • Off—The port does not use flow control, regardless of whether the neighboring port requests flow control. • On—The port uses flow control, as dictated by the neighboring port. • Desired—The port allows, but does not require, flow control frames. Flow control frames (also called pause frames) are special packets that signal a source to stop sending frames for a defined interval when buffers are full.
Flow Control Send	The flow control setting for outgoing frames: • Off—The port does not send flow control frames to the neighboring port. • On—The port sends flow control frames to the neighboring port. • Desired—The port allows, but does not require, flow control frames.
Roles	Lists the interface roles associated with the interface. Interface roles are objects that are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. See Understanding Interface Role Objects, page 6-73.

Create and Edit Interface Dialog Boxes—Dynamic Mode

Use the Create Interface dialog box (or the Edit Interface dialog box) to configure the attributes of physical and virtual interfaces that run in dynamic mode. Dynamic ports can convert the link into a trunk link based on the settings of the neighboring port.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, click **Add** or **Edit** to open the Create/Edit Interface dialog box, then select **Dynamic** from the Mode list.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12
- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14

- Interface Auto Name Generator Dialog Box, page 62-12
- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 68-7 Create and Edit Interface Dialog Boxes—Dynamic Mode

Element	Description
Enable Interface	When selected, enables the interface. When deselected, disables the interface using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface. For details about defining a subinterface, see Create and Edit Interface Dialog Boxes—Subinterfaces, page 68-22.
Name (Select button)	Displays the generated interface name, if the name has been set. Click **Select** to open the Interface Auto Name Generator Dialog Box, page 62-12. From here, you can enter or edit the details that Security Manager uses to generate an interface name.
Mode	The port configuration type for this interface. Select **Dynamic** to display the configuration options that are relevant for dynamic ports.
Dynamic Port settings	
Dynamic Mode	The dynamic trunk mode: • Auto—Allows the port to convert the link to a trunk link. The port becomes a trunk port if the neighboring port is set to Trunk or Desirable mode. • Desirable—Makes the port actively attempt to convert the link to a trunk link.
Access VLAN ID	The access VLAN ID to use when the port does *not* function as a trunking link. This can occur when the neighboring interface is not set to trunk, auto, or desirable mode. Valid values range from 1 to 4094.

Table 68-7 Create and Edit Interface Dialog Boxes—Dynamic Mode (continued)

Element	Description
Encapsulation	Select one of the following: • DOT1Q—Specifies VLAN encapsulation on the trunk link, as defined by the IEEE 802.1Q standard. Applies only to Ethernet subinterfaces. • ISL—Specifies ISL encapsulation on the trunk link. 10-Gigabit Ethernet ports do not support ISL encapsulation. • Negotiate—Specifies that the interface negotiates with the neighboring interface to become either an ISL or 802.1Q trunk, based on the configuration and capabilities of the neighboring interface. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the router.
Native VLAN (Select button)	Enables you to select the Native VLAN to associate with this interface, using the ID specified in the VLAN ID field. (If no VLAN ID is specified for the Native VLAN, the default is 1.) This option applies to you only if you are configuring a physical interface that is meant to serve as an 802.1Q trunk interface. You must first specify DOT1Q as the encapsulation type. The Native VLAN of a trunk interface is the VLAN to which all untagged VLAN packets are logically assigned. This includes the management traffic associated with the VLAN. When deselected, the Native VLAN is not associated with this interface. **Note** The Native VLAN cannot be configured on a subinterface of the trunk interface. Be sure to configure the same Native VLAN value at both ends of the link; otherwise, traffic may be lost or sent to the wrong VLAN. Click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can associate a native VLAN with the described interface.
Allowed VLANs (Select button)	Enables you to specify which VLANs are allowed on the trunk. Enter the VLAN IDs. Use commas to separate multiple VLANs or use a hyphen to indicate a range of VLANs (for example, 12,17,22 or 2-200). Valid IDs range from 1 to 4094. Alternatively, click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select the VLANs to include on the trunk.

Table 68-7 Create and Edit Interface Dialog Boxes—Dynamic Mode (continued)

Element	Description
Prune VLANs (Select button)	Enables you to specify which VLANs are eligible for pruning. Enter the VLAN IDs. Use commas to separate multiple VLANs or use a hyphen to indicate a range of VLANs (for example, 12,17,22 or 2-200.) Alternatively, click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select the VLANs that are eligible for pruning.
Enable VACL Capture	When selected, enables VACL capture. If the capture bit is set, ports with the capture function enabled can receive forwarded packets. When deselected, disables VACL capture.
Capture VLANs (Select button)	Enables you to identify the VLANs where VACLs should receive forwarded VLAN packets. This option is available if you selected the Enable VACL Capture check box. Enter a comma-separated list of VLAN IDs or click **Select** to open the VLAN Selector Dialog Box, page 68-35. VACLs can capture VLAN packets only when they are initially routed or bridged into the VLAN. Only forwarded packets can be captured.
Common interface settings	
Speed	The speed of the physical interface: • 10—Transmits at 10 Mbps. • 100—Transmits at 100 Mbps. • 1000—Transmits at 1,000 Mbps. • 10000—Transmits at 10,000 Mbps. • Auto—If Speed is set to Auto, both Speed and Duplex are autonegotiated. • Non-Negotiate—Disables link negotiation.
Duplex	The duplex setting of the interface: • Auto—Autonegotiates the duplex. • Half—Sends and receives data, but not at the same time • Full—Sends and receives data at the same time. If the speed is set to Auto, the duplex setting must also be set to Auto.
MTU	The maximum transmission unit, which refers to the largest packet size (in bytes) that can be handled by the interface. The range of valid values depends on the interface type.
Description	A text description of the interface. Enter up to 240 characters on a single line, without using carriage returns. **Note** For multiple context mode, the system description is independent of the context description.

Table 68-7 Create and Edit Interface Dialog Boxes—Dynamic Mode (continued)

Element	Description
Flow Control Receive	The flow control setting for incoming frames: • Off—The port does not use flow control, regardless of whether the neighboring port requests flow control. • On—The port uses flow control, as dictated by the neighboring port. • Desired—The port allows, but does not require, flow control frames. Flow control frames (also called pause frames) are special packets that signal a source to stop sending frames for a defined interval when buffers are full.
Flow Control Send	The flow control setting for outgoing frames: • Off—The port does not send flow control frames to the neighboring port. • On—The port sends flow control frames to the neighboring port. • Desired—The port allows, but does not require, flow control frames.
Roles	Lists the interface roles associated with the interface. Interface roles are objects that are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. See Understanding Interface Role Objects, page 6-73.

Create and Edit Interface Dialog Boxes—Subinterfaces

Use the Create Interface dialog box (or the Edit Interface dialog box) to configure the attributes of subinterfaces defined on Catalyst 6500/7600 devices.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, click **Add** or **Edit** to open the Create/Edit Interface dialog box, then select **Subinterface** from the Type list.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12
- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14
- Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18
- Understanding Interface Role Objects, page 6-73

Field Reference

Table 68-8 Create and Edit Interface Dialog Boxes—Subinterfaces

Element	Description
Enable Interface	When selected, enables the subinterface. When deselected, disables the subinterface using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface. Select **Subinterface**.
Parent	Identifies the parent interface of the subinterface.
Subint. ID	Specifies the ID for the subinterface. The numeric ID string cannot exceed 10 characters.
IP Type	The type of IP address used by the subinterface: • Static IP—Specifies that the subinterface uses a permanent IP address and activates related GUI elements.
IP Address	Enables you to enter an IP address.
Helper IP Addresses	Enables you to assign a helper IP address to the subinterface. A helper IP address converts broadcast DHCP requests to unicast requests that are directed exclusively to the DHCP server.
Mask	Enables you to specify the subnet mask. You can enter a netmask value or you can select a netmask from the list. If you enter a netmask, you can express its value in dotted decimal format (for example, 255.255.255.0) or you can enter the number of bits (for example, 24). **Note** Do not use 255.255.255.254 or 255.255.255.255 for any interface that is connected to your network; these netmasks cause all traffic on an interface to stop.
Encapsulation	The encapsulation type defined for the subinterface: • [blank]—No encapsulation is defined. • DOT1Q—Specifies VLAN encapsulation on the trunk link, as defined by the IEEE 802.1Q standard. Applies only to Ethernet subinterfaces. • ISL—Specifies ISL encapsulation on the trunk link. 10-Gigabit Ethernet ports do not support ISL encapsulation. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the router.
VLAN ID	Applies only when encapsulation is defined for the subinterface. The VLAN ID associated with the subinterface.
Description	A text description of the interface. Enter up to 240 characters on a single line, without using carriage returns. **Note** For multiple context mode, the system description is independent of the context description.

Create and Edit Interface Dialog Boxes—Unsupported Mode

If you discover an interface configured with a mode that is not supported by Security Manager (such as dot1q-tunnel or private-vlan), the interface is displayed in Unsupported mode. You can view the attributes of this interface, but you cannot make any changes to the configuration unless you first change the mode. All definition fields, other than Mode, are read-only.

Navigation Path

Go to the Interfaces/VLANs Page—Interfaces Tab, page 68-8, select an interface whose mode is defined as Unsupported, then click **Add** or **Edit** to open the Create/Edit Interface dialog box.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Create and Edit Interface Dialog Boxes—Routed Port Mode, page 68-12
- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14
- Create and Edit Interface Dialog Boxes—Dynamic Mode, page 68-18

Field Reference

Table 68-9 Create and Edit Interface Dialog Boxes—Unsupported Mode

Element	Description
Enable Interface	When selected, indicates that the interface is enabled. When deselected, indicates that the interface has been disabled using the shutdown command.
Type	Specifies whether the definitions apply to an interface or a subinterface.
Name (Select button)	Displays the name of the interface.
Mode	Displays Unsupported, which designates an interface whose mode is not supported by Security Manager. Select a different option to change the interface mode. **Note** If you change the interface mode, you can then modify the other settings in this dialog box.
Speed	Displays the speed of the physical interface: • 10—Transmits at 10 Mbps. • 100—Transmits at 100 Mbps. • 1000—Transmits at 1,000 Mbps. • 10000—Transmits at 10,000 Mbps. • Auto—If Speed is set to Auto, both Speed and Duplex are autonegotiated. • Non-Negotiate—Disables link negotiation.

Table 68-9 Create and Edit Interface Dialog Boxes—Unsupported Mode (continued)

Element	Description
Duplex	Displays the duplex setting of the interface: • Auto—Autonegotiates the duplex. • Half—Sends and receives data, but not at the same time • Full—Sends and receives data at the same time. If the speed is set to Auto, the duplex setting must also be set to Auto.
MTU	Displays the maximum transmission unit, which refers to the largest packet size (in bytes) that can be handled by the interface. The range of valid values depends on the interface type.
Description	Displays a text description of the interface. For multiple context mode, the system description is independent of the context description.
Flow Control Receive	Displays the flow control setting for incoming frames: • Off—The port does not use flow control, regardless of whether the neighboring port requests flow control. • On—The port uses flow control, as dictated by the neighboring port. • Desired—The port allows, but does not require, flow control frames. Flow control frames (also called pause frames) are special packets that signal a source to stop sending frames for a defined interval when buffers are full.
Flow Control Send	Displays the flow control setting for outgoing frames: • Off—The port does not send flow control frames to the neighboring port. • On—The port sends flow control frames to the neighboring port. • Desired—The port allows, but does not require, flow control frames.
Roles	Lists the interface roles associated with the interface. Interface roles are objects that are replaced with the actual interface IP addresses when the configuration is generated for each device. They allow you to define generic rules—ones that can apply to multiple interfaces. See Understanding Interface Role Objects, page 6-73.

VLANs

A VLAN is a switched network that is segmented logically instead of on the basis of geography. For example, a VLAN might interconnect members of a geographically dispersed workgroup. VLANs offer a practical convenience for many organizations because they reduce the need to rearrange the physical placement of personnel, equipment, and network infrastructure. Properly configured VLANs are scalable, secure, and can simplify the tasks of network management.

A VLAN consists of hosts and network devices (such as bridges and routers), connected by a single bridging domain. Traffic between VLANs must be routed.

Security Manager helps you to create VLANs and define VLAN settings for the defined interfaces on Cisco Catalyst switches and Cisco 7600 Series routers, their supported services modules, and their security contexts.

The following topics describe the actions you can perform when defining VLANs on Catalyst devices:

- Creating or Editing VLANs, page 68-26
- Deleting VLANs, page 68-27
- Interfaces/VLANs Page—VLANs Tab, page 68-27

Related Topics

- VLAN Groups, page 68-31
- VLAN ACLs (VACLs), page 68-36
- Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers"

Creating or Editing VLANs

You can create a VLAN or reconfigure the attributes of a VLAN.

Related Topics

- Deleting VLANs, page 68-27
- Creating or Editing VLAN Groups, page 68-32
- Creating or Editing VACLs, page 68-37
- Create and Edit VLAN Dialog Boxes, page 68-28
- VLANs, page 68-25

Step 1 (Device view) Select a Catalyst device, select Interfaces/VLANs from the Policy selector, then click the VLANs tab in the work area.

The VLANs tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—VLANs Tab, page 68-27.

Step 2 Do one of the following:

- To define the attributes of a new VLAN, click **Add Row**.
- To edit the attributes of a VLAN, select it in the list, then click **Edit Row**.

See Create and Edit VLAN Dialog Boxes, page 68-28, for a description of the fields in the dialog box.

Step 3 In the VLAN ID field, enter a unique ID number for the VLAN. The number that you enter must not be assigned to any other VLAN in the bridging group.

Step 4 (Optional) Enter a name for the VLAN.

Step 5 (Optional) If the VLAN is part of a VLAN group, select the group ID, or select **Add Group** to open the Create VLAN Group dialog box. For more information, see Creating or Editing VLAN Groups, page 68-32.

Step 6 From the Status list, specify the status of the VLAN (active or suspended).

Step 7 From the Type list, select either **Layer 2** or **Layer 3**.

Step 8 (Optional) For a Layer 3 VLAN, define a switched virtual interface (SVI):

a. To make the SVI active, select the **Enable Interface** check box. An SVI enables routing between VLANs and provides IP host connectivity to the switch. If you do not select this check box, the SVI is created in shutdown mode.

b. Enter the IP address for the SVI.

c. Enter the SVI subnet mask by typing it, or select a netmask value from the Subnet Mask list.

d. Enter an optional description, if required.

Step 9 Do one or both of the following:

- To associate access ports with the VLAN, enter their names in the Access Ports text box or click **Select** to open an interface selector.
- To associate trunk ports with the VLAN, enter their names in the Trunk Ports text box or click **Select** to open an interface selector.

See Interface Selector Dialog Box—VLAN ACL Content, page 68-42 for a description of the fields in the dialog box. For more information about defining ports, see Creating or Editing Ports on Cisco Catalyst Switches and Cisco 7600 Series Routers, page 68-6.

Step 10 Click **OK** to save your definitions locally on the client and close the dialog box.

Deleting VLANs

You can delete a VLAN. However, deleting a VLAN does not delete it from any policy that might reference it. Ensure that your other policies do not use the VLAN before you delete it. When you submit your changes to the database, Security Manager points out any undefined VLANs that are referenced by other policies.

Related Topics

- Creating or Editing VLANs, page 68-26
- VLANs, page 68-25

Step 1 (Device view) Select a Cisco Catalyst switch or Cisco 7600 Series router from the Device selector.

Step 2 Select **Interfaces/VLANs** from the Policies selector.

Step 3 Click the VLANs tab in the work area.

The VLANs tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—VLANs Tab, page 68-27.

Step 4 Select a VLAN from the table, then click **Delete Row**.

The VLAN is deleted.

Interfaces/VLANs Page—VLANs Tab

Use the VLANs tab to view and configure VLANs on supported Cisco Catalyst switches and Cisco 7600 Series routers.

Navigation Path

- (Device view) Select **Interfaces/VLANs** from the Device selector, then click the **VLANs** tab.

Related Topics

- Interfaces/VLANs Page—VLAN Groups Tab, page 68-33
- Interfaces/VLANs Page—Interfaces Tab, page 68-8
- Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups, page 68-3
- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Create and Edit VLAN Dialog Boxes, page 68-28
- Filtering Tables, page 1-48

Field Reference

Table 68-10 Interfaces/VLANs Page—VLANs Tab

Element	Description
VLAN ID	Interface-specific identity of the VLAN that a table row describes. The VLAN ID specifies where 802.1Q tagged packets are sent and received on the subinterface; without a VLAN ID, the subinterface cannot send or receive traffic. Valid values range from 2 to 4094 (VLAN ID 1 is reserved). **Note** All VLAN IDs must be unique among all subinterfaces configured on the same physical interface. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the device.
Name	Name of the corresponding VLAN for an interface or subinterface.
Interface	Identifies the logical name of the interface (interface role) or physical interface.
Type	Specifies whether a VLAN has access to Layer 2 or Layer 3.
Status	Indicates whether a VLAN is active or suspended.
Add Row button	Opens the Create VLAN dialog box to define a new VLAN.
Edit Row button	Opens the Edit VLAN dialog box to edit the selected VLAN.
Delete Row button	Deletes the selected VLAN.

Create and Edit VLAN Dialog Boxes

Use the Create VLAN dialog box (or the Edit VLAN dialog box) to configure or reconfigure VLAN settings and attributes.

Navigation Path

Go to the Interfaces/VLANs Page—VLANs Tab, page 68-27, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Understanding FlexConfig Policies and Policy Objects, page 7-2
- Create and Edit VLAN Group Dialog Boxes, page 68-34
- Interface Selector Dialog Box—VLAN ACL Content, page 68-42

Field Reference

Table 68-11 Create and Edit VLAN Dialog Box

Element	Description
VLAN ID	Displays the VLAN ID if one is configured. Otherwise, enter the ID manually. The VLAN ID specifies where 802.1Q tagged packets are sent and received on an interface or subinterface; without a VLAN ID, the interface or subinterface cannot send or receive traffic. Each VLAN must have an ID. Valid values range from 1 to 4094. **Note** All VLAN IDs must be unique among all subinterfaces configured on the same physical interface. **Tip** To configure DOT1Q encapsulation on an Ethernet interface without associating the VLAN with a subinterface, enter the **vlan-id dot1q** command using CLI commands or FlexConfigs. Configuring VLANs on the main interface increases the number of VLANs that can be configured on the device.
Name	Enter a name for the VLAN, or view the VLAN name if you entered one previously. Each VLAN must have an ID, and can optionally have a name. The maximum length is 32 characters.
Group	The VLAN group to which the VLAN belongs. A VLAN can be associated with one group only. You can associate the VLAN with an existing group, or select **Add Group** to open the Create VLAN Group dialog box.
Status	The current status of the VLAN: • Active—The VLAN carries traffic. • Suspended—The VLAN does not pass packets.
Type	Indicates whether the specified VLAN is configured for Layer 2 or Layer 3, and enables you to choose the kind of VLAN that you prefer. A Layer 3 VLAN requires an IP address and creates a VLAN interface.

Table 68-11 Create and Edit VLAN Dialog Box (continued)

Element	Description
Switch Virtual Interface	Applies only when defining a Layer 3 VLAN. • Enable Interface—When selected, enables the switched virtual interface (SVI), which is a virtual interface that you can attach to any VLAN. The SVI enables routing between VLANs and provides IP host connectivity to the switch. When deselected, disables the SVI. • IP Address—The IP address for the SVI. An IP address is required for management access. • Subnet Mask—The subnet mask for the SVI. Select any option from the list of valid subnet mask entries. • Description—Enables you to enter a description of up to 240 characters on a single line, without carriage returns. For multiple context mode, the system description is independent of the context description.
Access Ports (Select button)	Lists which access ports are associated with the specified VLAN, if any are associated, and enables you to add or remove access port associations for the specified VLAN. You can associate any number of access ports with a VLAN. Click **Select** to open the Access Port Selector Dialog Box, page 68-30. From here, you can associate access ports with the specified VLAN, or remove access port associations from the VLAN.
Trunk Ports (Select button)	Lists which trunk ports are associated with the specified VLAN, if any are associated, and enables you to add or remove trunk port associations for the specified VLAN. A VLAN can belong to the allowed list of one or more trunk ports. You can include a VLAN in a trunk port group. Click **Select** to open the Trunk Port Selector Dialog Box, page 68-31. From here, you can associate trunk ports with the specified VLAN, or remove trunk port associations from the VLAN.

Access Port Selector Dialog Box

Use the Access Port Selector dialog box to define which access ports are associated with a selected VLAN.

Navigation Path

Open the Create and Edit VLAN Dialog Boxes, page 68-28, then click **Select** in the Access Ports field.

Related Topics

- Create and Edit Interface Dialog Boxes—Access Port Mode, page 68-9
- Trunk Port Selector Dialog Box, page 68-31
- Filtering Tables, page 1-48

Field Reference

Table 68-12 Access Port Selector Dialog Box

Element	Description
Available Access Ports	Displays the access ports that are not assigned to a particular VLAN.
Add >> button	Adds interfaces that are selected in the Available Access Ports list to the Selected Access Ports list.
Remove << button	Removes selected interfaces from the Selected Access Ports list.
Selected Access Ports	Displays the interface objects that are selected.
Add Row button	Opens the Create Interface dialog box to define a new interface.
Edit Row button	Opens the Edit Interface dialog box to edit the selected interface.

Trunk Port Selector Dialog Box

Use the Trunk Port Selector dialog box to define which trunk ports are associated with a selected VLAN.

Navigation Path

Open the Create and Edit VLAN Dialog Boxes, page 68-28, then click **Select** in the Trunk Ports field.

Related Topics

- Create and Edit Interface Dialog Boxes—Trunk Port Mode, page 68-14
- Access Port Selector Dialog Box, page 68-30
- Filtering Tables, page 1-48

Field Reference

Table 68-13 Trunk Port Selector Dialog Box

Element	Description
Available Trunk Ports	Displays all available trunk ports.
Add >> button	Adds interfaces that are selected in the Available Trunk Ports list to the Selected Trunk Ports list.
Remove << button	Removes selected interfaces from the Selected Trunk Ports list.
Selected Trunk Ports	Displays the interface objects that are selected.
Add Row button	Opens the Create Interface dialog box to define a new interface.
Edit Row button	Opens the Edit Interface dialog box to edit the selected interface.

VLAN Groups

A VLAN group defines a logical collection of VLANs. The VLAN Groups tab on the Interfaces/VLANs page displays:

- All VLAN groups that are defined on the selected device.
- The service module slots to which a VLAN group is bound.

- Which VLANs belong to each VLAN group.

VLAN groups can be used when assigning VLANs to an FWSM security context. A VLAN group can be assigned to multiple FWSMs, and each FWSM can have multiple VLAN groups assigned to it. To perform this assignment, see Add/Edit Security Context Dialog Box (FWSM), page 59-8.

The following topics describe the actions you can perform when defining VLAN groups on Catalyst devices:

- Creating or Editing VLAN Groups, page 68-32
- Deleting VLAN Groups, page 68-33
- Interfaces/VLANs Page—VLAN Groups Tab, page 68-33

Related Topics

- Interfaces, page 68-5
- VLANs, page 68-25
- VLAN ACLs (VACLs), page 68-36
- Chapter 68, “Managing Cisco Catalyst Switches and Cisco 7600 Series Routers”

Creating or Editing VLAN Groups

You can create VLAN groups. When you create a VLAN group, remember that:

- Each group must have an ID.
- You can associate a VLAN group with one or more FWSM modules.
- Each VLAN can be a member of only one VLAN group.

Related Topics

- Deleting VLAN Groups, page 68-33
- Creating or Editing VLANs, page 68-26
- Creating or Editing VACLs, page 68-37
- Interfaces/VLANs Page—VLAN Groups Tab, page 68-33
- VLAN Groups, page 68-31

Step 1 (Device view) Select a Catalyst device, select **Interfaces/VLANs** from the Policy selector, then click the VLAN Groups tab in the work area.

The VLAN Groups tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—VLAN Groups Tab, page 68-33.

Step 2 Do one of the following:

- To define the attributes of a new VLAN group, click **Add Row**.
- To edit the attributes of a VLAN group, select it in the list, then click **Edit Row**.

See Create and Edit VLAN Group Dialog Boxes, page 68-34, for a description of the fields in this dialog box.

Step 3 In the VLAN Group ID field, enter a unique ID number for the VLAN group. The number that you enter must not be assigned to any other VLAN group.

Step 4 To associate the VLAN group with specific service module slots, enter their slot numbers in the Service Module Slots text box, or click **Select** to open a selector.

Note Defining this association makes it possible to later assign this VLAN group to a security context on the FWSM. See Add/Edit Security Context Dialog Box (FWSM), page 59-8.

Step 5 Enter the VLANs to add to the VLAN group, or click **Select** to open a selector.

Step 6 Click **OK** to save your definitions locally on the client and close the dialog box.

Deleting VLAN Groups

You can delete VLAN groups. Deleting a VLAN group has no affect on the VLANs in the group.

Related Topics

- Creating or Editing VLAN Groups, page 68-32
- VLAN Groups, page 68-31

Step 1 (Device view) Select a Catalyst 6500 Series switch or Cisco 7600 Series router from the Device selector.

Step 2 Select **Interfaces/VLANs** from the Policy selector.

Step 3 Click the VLAN Groups tab in the work area.

The VLANs tab is displayed. For a description of the fields on this tab, see Interfaces/VLANs Page—VLAN Groups Tab, page 68-33.

Step 4 Select a VLAN group from the table, then click **Delete Row**. The VLAN group is deleted.

Interfaces/VLANs Page—VLAN Groups Tab

Use the VLAN Groups tab to view and configure VLAN groups on supported 6500 Series switches and 7600 Series routers.

Note The VLAN Groups tab is available only for Cisco Catalyst 6500 Series switches and Cisco 7600 Series routers.

Navigation Path

- (Device view) Select **Interfaces/VLANs** from the Device selector, then click the **VLAN Groups** tab.

Related Topics

- Interfaces/VLANs Page—VLANs Tab, page 68-27
- Interfaces/VLANs Page—Interfaces Tab, page 68-8
- Viewing a Summary of Catalyst Interfaces, VLANs, and VLAN Groups, page 68-3

- Create and Edit VLAN Group Dialog Boxes, page 68-34
- Filtering Tables, page 1-48

Field Reference

Table 68-14 Interfaces/VLANs Page—VLAN Groups Tab

Element	Description
VLAN Group	Numeric ID of a VLAN group that is configured on the selected device.
Service Module Slots	Associates the chassis slot number (in which the relevant services module is installed) with the interface through which a particular VLAN participates in the VLAN group.
VLAN IDs	The VLAN IDs associated with this group. Valid values range from 1 to 65535.
Add Row button	Opens the Create VLAN Group dialog box to define a new VLAN group.
Edit Row button	Opens the Edit VLAN Group dialog box to edit the selected VLAN group.
Delete Row button	Deletes the selected VLAN group.

Create and Edit VLAN Group Dialog Boxes

Use the Create and Edit VLAN Group dialog box to configure or reconfigure the attributes of VLAN groups, which are logical groups of VLANs that you want to associate with one another when you define VLAN port policies.

Navigation Path

Do one of the following:

- Go to the Interfaces/VLANs Page—VLAN Groups Tab, page 68-33, then click the **Add** or **Edit** button beneath the table.
- Go to the Interfaces/VLANs Page—VLANs Tab, page 68-27, click the **Add** or **Edit** button beneath the table, then select **Add Group** from the Group list.

Related Topics

- Service Module Slot Selector Dialog Box, page 68-35

Field Reference

Table 68-15 Create and Edit VLAN Group Dialog Boxes

Element	Description
VLAN Group ID	The 802.1q VLAN group name. Valid values range from 1 to 65535.

Table 68-15 Create and Edit VLAN Group Dialog Boxes (continued)

Element	Description
Service Module Slots (Select button)	The chassis slot number (in which the relevant services module is installed) that is associated with the interface through which a particular VLAN participates in the VLAN group. Enter the slot number or click **Select** to open the Service Module Slot Selector Dialog Box, page 68-35. **Note** After you associate the VLAN group with a service module, such as an FWSM, you can assign the VLAN group to the security contexts of the FWSM. See Add/Edit Security Context Dialog Box (FWSM), page 59-8.
VLAN IDs (Select button)	The comma-separated IDs of all VLANs that are part of the group. Each VLAN can be a member of only one group. Click **Select** to open the Service Module Slot Selector Dialog Box, page 68-35. From here, you can select VLANs to include in the VLAN group.

Service Module Slot Selector Dialog Box

Use the Service Module Slot Selector dialog box to associate a service module with a VLAN.

Navigation Path

Go to the Create and Edit VLAN Group Dialog Boxes, page 68-34, then click **Select** in the Service Module Slots field.

Related Topics

- VLAN Selector Dialog Box, page 68-35
- Filtering Tables, page 1-48

Field Reference

Table 68-16 Service Module Selector Dialog Box

Element	Description
Available Service Module Slots	Displays the defined service module slots.
Add >> button	Moves selected service module slots from the Available Service Module Slots list to the Selected Service Module Slots list.
Remove << button	Removes selected service module slots from the Selected Service Modules list.
Selected Service Module Slots	Displays the selected service module slots.

VLAN Selector Dialog Box

Use the VLAN Selector dialog box to associate VLANs with interfaces, VLAN groups, security contexts, and VACLs.

Navigation Path

You can access this dialog box when you define interfaces, VLAN groups, IDSM settings, or VACLs by clicking the **Select** button in any field used for defining VLANs.

Related Topics

- Service Module Slot Selector Dialog Box, page 68-35
- Filtering Tables, page 1-48

Field Reference

Table 68-17 VLAN Selector Dialog Box

Element	Description
Available VLANs	Displays defined VLANs that are available to be associated with the object you are configuring. **Note** The VLANs that are available will depend on the type of object you are configuring and other settings defined on the device. For example, when selecting VLANs to assign to a VLAN group, the Available VLANs list will only contain VLANs that have not been assigned to another VLAN group. When selecting VLANs to assign to a security context, the Available VLANs list will only contain VLANs that are part of a VLAN group that has been assigned to the service module you are configuring.
Add >> button	Moves selected VLANs from the Available VLANs list to the Selected VLANs list.
Remove << button	Removes selected VLANs from the Selected VLANs list.
Selected VLANs	Displays the selected VLANs.
VLAN Ranges	The VLAN ranges entered manually before the selector was opened, if any.

VLAN ACLs (VACLs)

Cisco IOS standard or extended ACLs are configured on router interfaces only, and are applied on routed packets only. In contrast, Cisco Catalyst switches and Cisco 7600 Series routers can use VLAN ACLs (VACLs) to control the access of all packets that are bridged within a VLAN or that are routed to or from a VLAN for VACL capture through a WAN interface. VACLs:

- Are processed in hardware.
- Use Cisco IOS ACLs.
- Ignore any Cisco IOS ACL fields that are not supported in hardware.

Note Security Manager does not support the creation or configuration of MAC ACLs (MACLs), which are named ACLs that are sometimes used with VACLs to filter IPX, DECnet, AppleTalk, VINES, or XNS traffic based on MAC addresses.

When you configure a VACL and apply it to a VLAN, all packets entering the VLAN are checked against the VACL.

If you apply a VACL to a VLAN and you apply an ACL to a routed interface in that same VLAN, any packet coming into the VLAN is first checked against the VACL. Then, if permitted, the packet is checked against the input ACL before it reaches the routed interface.

When a packet is routed from one VLAN to another, it is first checked against the output ACL that is applied to the routed interface. Then, if permitted, the packet is checked against any VACLs that are configured for the destination VLAN.

If a VACL is configured for a packet type, and a packet of that type does not match the VACL, the default action is deny.

VLAN Access Maps

Security Manager uses *VLAN access maps* to configure VACLs. Conceptually similar to a route map, a VLAN access map is a container in which you place one or more *statements* (conditions that match an action) and number them by their order of importance. A VLAN access map must also identify the VLANs to which it is applied, contain the map name, and identify at least one VACL sequence.

A VACL sequence must have a sequence number and at least one action, and must match at least one ACL.

Devices evaluate map statements in sequence and you can associate more than one VLAN access map with any device chassis.

To manage a VACL, select a Catalyst device in Device View, then select **Platform > VLAN Access Lists**. You use VLAN access maps to configure VACLs for IP traffic.

The following topics describe the actions you can perform when defining VACLs on Catalyst devices:

- Creating or Editing VACLs, page 68-37
- Deleting VACLs, page 68-38
- VLAN Access Lists Page, page 68-39

Related Topics

- VLANs, page 68-25
- VLAN Groups, page 68-31
- Chapter 68, “Managing Cisco Catalyst Switches and Cisco 7600 Series Routers”

Creating or Editing VACLs

When you can create or edit a VACL, you must:

- Name the VACL.
- Define the VLANs to which the VACL applies.
- Define a sequence map containing at least one VACL sequence.

Related Topics

- Deleting VACLs, page 68-38
- Creating or Editing VLANs, page 68-26
- Creating or Editing VLAN Groups, page 68-32
- Create and Edit VLAN ACL Dialog Boxes, page 68-41

- VLAN Access Lists Page, page 68-39

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > VLAN Access Lists** from the Policy selector.
- (Policy view) Select **Catalyst Platform > VLAN Access Lists**.

The VLAN Access Lists page is displayed. For a description of the fields on this page, see VLAN Access Lists Page, page 68-39.

Step 2 Do one of the following:

- To define the attributes of a new VACL, click **Add Row**.
- To edit the attributes of a VACL, select it in the list, then click **Edit Row**.

A dialog box opens. See Create and Edit VLAN ACL Dialog Boxes, page 68-41, for a description of the fields in the dialog box.

Step 3 Enter a name for the VACL in the **VLAN ACL Name** field.

Step 4 In the VLANs field, specify the VLANs to which the VACL should be applied, or click **Select** to open a VLAN selector.

Step 5 Define the sequence map:

a. Click **Add Row** or **Edit Row** beneath the Sequence Map table. A dialog box opens. See Create and Edit VLAN ACL Content Dialog Boxes, page 68-41.

b. Enter a number to identify the sequence.

c. Specify the standard and extended ACLs to assign to the sequence, or click **Select** to select the ACL object from a list or to create a new ACL object. For more information about ACL objects, see Creating Access Control List Objects, page 6-53.

d. Specify the action to perform on traffic that matches the ACLs defined in this sequence. (When you select Redirect as the action, you must specify the physical destination interfaces, or click **Select** to display a selector. See Specifying Interfaces During Policy Definition, page 6-76.)

e. Click **OK** to save your definitions locally on the client and close the dialog box. The sequence is displayed in the Sequence Map table.

f. Repeat the process to add sequences to the sequence map.

g. Use the up and down arrows to reorder the sequences, if required.

Note The order in which you place the sequences is significant. When a flow matches a permit ACL entry, the associated action is taken without checking the remaining sequences. When a flow matches a deny ACL entry, it is checked against the next ACL in the same sequence or the next sequence. If a flow does not match any ACL entry and at least one ACL is configured for that packet type, the packet is denied.

Deleting VACLs

You can delete a VACL if it is not being used by any device, policy, or object.

Before You Begin

You must delete all references to the VACL before you can remove it from the database. To locate all references to the VACL, run an object usage report for it. See Generating Object Usage Reports, page 6-15.

Related Topics

- Creating or Editing VACLs, page 68-37
- Interfaces/VLANs Page—VLANs Tab, page 68-27
- VLAN ACLs (VACLs), page 68-36

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > VLAN Access Lists** from the Policy selector.
- (Policy view) Select **Catalyst Platform > VLAN Access Lists**.

The VLAN Access Lists page is displayed. For a description of the fields on this page, see VLAN Access Lists Page, page 68-39.

Step 2 Click in a row to select a VACL, then click **Delete**.

Step 3 Click **OK** to save your changes.

VLAN Access Lists Page

Use the VLAN Access Lists page to view and configure VLAN access lists for Cisco Catalyst switches and Cisco 7600 Series routers.

Navigation Path

You can access this page from:

- (Device view) Select **Platform > VLAN Access Lists** from the Device Policy selector.
- (Device view) Select **Catalyst Platform >VLAN Access Lists** from the Policy Types selector.

Related Topics

- Creating Access Control List Objects, page 6-53
- Create and Edit VLAN ACL Dialog Boxes, page 68-41
- Create and Edit VLAN ACL Content Dialog Boxes, page 68-41
- Filtering Tables, page 1-48

Field Reference

Table 68-18 VLAN Access Lists Page

Element	Description
VLAN Access Lists table	
VLAN ACL	Displays the VLAN ACL name.

Table 68-18 VLAN Access Lists Page (continued)

Element	Description
Sequence	Specifies the map sequence number. VACL sequences are applied in order of sequence, from lowest number to highest.
Matching	Displays the Match ACLs, if any are defined. VACL matching occurs only when an ACL permit is encountered. ACL denies are ignored.
Action	Specify whether the action is to drop, drop and log, forward, forward and capture, or redirect packets. **Note** The redirect action helps you to specify as many as five interfaces, which can be physical interfaces or EtherChannels. You cannot redirect packets to an EtherChannel member or a VLAN interface.
VLAN IDs	Interface-specific identity of the VLAN that a table row describes. The VLAN ID specifies where 802.1Q tagged packets are sent and received on the subinterface; without a VLAN ID, the subinterface cannot send or receive traffic.
Add Row button	Opens the Create VLAN ACL dialog box, where you can define a new VACL.
Edit Row button	Opens the Edit VLAN ACL dialog box, where you can edit the selected VACL.
Delete Row button	Deletes the selected access list.
Additional fields	
Log Table Size	Displays the log table size. Valid sizes range from 0 to 2048 and the default is 500. Logged packets from new flows are dropped when the table is full.
Max. Packet Rate	Displays the maximum redirect VACL logging packet rate per second. Valid rates range from 10 to 5000 packets per second and the default rate is 2000. Packets that exceed the limit are dropped.
Logging Threshold	Displays the logging threshold if one is set. By default, no threshold is set. When you configure VACL logging, IP packets that are denied generate log messages on a per-flow basis if the threshold for a flow is reached in any interval of less than 5 minutes. Only dropped IP packets can be logged.
Capture Interfaces	Identifies the interface that captures forwarded packets in which the capture bit is set. You can configure any interface as the capture interface. The capture action sets the capture bit for the forwarded packets so that ports with the capture function enabled can receive the packets. Only forwarded packets can be captured. **Note** The information shown here is read-only. To define capture interfaces, use the Create/Edit Interface dialog box. See Interfaces/VLANs Page—Interfaces Tab, page 68-8.

Create and Edit VLAN ACL Dialog Boxes

Use the Create VLAN ACL dialog box (or the Edit VLAN ACL dialog box) to configure or reconfigure VACL attributes.

Navigation Path

Go to the VLAN Access Lists Page, page 68-39, then click the **Add** or **Edit** button beneath the table.

Related Topics

- Create and Edit VLAN Dialog Boxes, page 68-28
- Create and Edit VLAN Group Dialog Boxes, page 68-34
- Filtering Tables, page 1-48

Field Reference

Table 68-19 Create and Edit VLAN ACL Dialog Boxes

Element	Description
VLAN ACL Name	The user-defined name for the VACL.
VLANs (Select button)	Enables you to designate the VLANs to which the VACL should be applied. Do one of the following: • Enter VLAN IDs. You can use commas to separate multiple VLANs or use a hyphen to indicate a range of VLANs. For example: 12,17,22 or 2-200. Valid IDs range from 1 to 4094. • Click **Select** to open the VLAN Selector Dialog Box, page 68-35.
Sequence Map table	The sequence maps included in the VLAN access map. A VLAN access map can consist of one or more map sequences, where each sequence pairs a *match clause*, which specifies an ACL object for traffic filtering, to an *action clause*, which specifies the action to take on packets that meet the criteria defined in the match ACLs. • To add a sequence map, click the Add Row (+) button and fill in the Create VLAN ACL Content dialog box (see Create and Edit VLAN ACL Content Dialog Boxes, page 68-41). • To edit a sequence map, select it and click the Edit Row button. • To delete a sequence map, select it and click the Delete Row button. • To change the order of a map, select it and click the Up or Down arrow buttons until it is in the desired position. The sequence number changes as you move it.

Create and Edit VLAN ACL Content Dialog Boxes

Use the Create VLAN ACL Content dialog box (or the Edit VLAN ACL Content dialog box) to configure or reconfigure VACL sequences.

Navigation Path

Go to the Create and Edit VLAN ACL Dialog Boxes, page 68-41, then click the **Add** or **Edit** button beneath the Sequence Map table.

Related Topics

- Create and Edit VLAN Dialog Boxes, page 68-28
- Create and Edit VLAN Group Dialog Boxes, page 68-34

Field Reference

Table 68-20 Create and Edit VLAN ACL Content Dialog Boxes

Element	Description
Sequence	Specify the map sequence number for the VLAN access map. Valid values range from 1 to 65535.
Match ACLs	Specify which ACLs the sequence should include in its match clause. Enter the names of the standard and extended ACL objects to include in the sequence, or click **Select** to select them from a list or to create new ones. You cannot use a MAC-layer ACL.
Action	The option to perform on packets that meet the criteria defined in the match ACLs: • Drop—Drops the packets. • Drop/Log—Logs the dropped packets. • Forward—Forwards the packets to their destination (using hardware switching). • Forward/Capture—Sets the capture bit for the forwarded packets so that ports with the capture function enabled also receive the packets. • Redirect—Redirects packets to the Ethernet interfaces defined in the Interfaces field.
Interfaces (Select button)	Applies only when the specified action is Redirect. The destination interfaces for redirect packets. Enter the names of up to five physical interfaces, or click **Select** to open the Interface Selector Dialog Box—VLAN ACL Content, page 68-42. The redirect interfaces must be in the VLAN for which the VACL access map is configured. **Note** You cannot redirect packets to an EtherChannel member or a VLAN interface. You also cannot redirect packets to a subinterface.

Interface Selector Dialog Box—VLAN ACL Content

Use the Interface Selector dialog box to define redirect interfaces when you create entries for a VACL sequence map.

Navigation Path

Open the Create and Edit VLAN ACL Content Dialog Boxes, page 68-41, select **Redirect** as the action, then click **Select** in the Interfaces field.

Related Topics

- Create and Edit VLAN ACL Dialog Boxes, page 68-41
- VLAN Access Lists Page, page 68-39
- Filtering Tables, page 1-48

Field Reference

Table 68-21 Interface Selector Dialog Box

Element	Description
Available Interfaces	Displays the physical interfaces that are defined in the Interfaces/VLANs policy.
Add >> button	Adds interfaces that are selected in the Available Interfaces list to the Selected Interfaces list.
Remove << button	Removes selected interfaces from the Selected Interfaces list.
Selected Interfaces	Displays the interfaces that are selected.

IDSM Settings

When you select a Catalyst device in Device view, then select **Platform > IDSM Settings** from the Policy selector, a list is displayed that:

- Displays the settings for data ports on Intrusion Detection System Service Modules (IDSMs).
- Helps you to organize IDSM data ports in channel groups.

The IDSM card detects and stops security threats on network connections. The card inspects the traffic that enters its two data ports and drops packets if a security threat is detected. The data port settings define:

- Which traffic is received by the data ports, as defined by the VLAN IDs.
- The sensing mode used by the data ports:
  - Trunk (IPS)—The IDSM performs VLAN bridging between pairs of VLANs within the same data port, operating as an 802.1q trunk. The IDSM inspects the traffic it receives on each VLAN in a VLAN pair and can either forward the packets on the other VLAN in the pair or drop the packet if an intrusion attempt is detected.
  - Capture (IDS)—The IDSM passively monitors network traffic that was copied to the data ports by the Catalyst switch using either VACL capture or SPAN. The data ports operate as 802.1q trunks that can be configured to trunk different VLANs. When operating in this passive mode, the IDSM cannot drop packets in response to a network intrusion attempt, but it can send TCP resets over the data ports in an attempt to block the intrusion.

Note Security Manager supports a subset of IDSM settings on chassis running IOS 12.2(18)SXF4 or later. Trunk (IPS) and Capture (IDS) modes are supported; inline mode is not supported. Security Manager cannot manage IDSM data ports that are part of a spanning tree or access VLAN.

For high-traffic networks, EtherChannel is used to perform load balancing among multiple data ports. These data ports might be located on different IDSM cards within the same Catalyst device.

EtherChannel is also used to redirect traffic in the event of port failure to the remaining ports within the channel group. This resiliency help preserve intrusion detection and prevention without user intervention and with minimum packet loss.

The following topics describe the actions you can perform when defining IDSM settings:

- Creating or Editing EtherChannel VLAN Definitions, page 68-44
- Deleting EtherChannel VLAN Definitions, page 68-45
- Creating or Editing Data Port VLAN Definitions, page 68-46
- Deleting Data Port VLAN Definitions, page 68-47
- IDSM Settings Page, page 68-47

Related Topics

- VLANs, page 68-25
- Chapter 68, "Managing Cisco Catalyst Switches and Cisco 7600 Series Routers"

Creating or Editing EtherChannel VLAN Definitions

When defining an EtherChannel VLAN definition, you must:

- Define the slot-port combination containing the data ports to include in the channel group.
- Select the sensing mode used by the data ports.
- Define which VLANs are forwarded to the data ports.

The following restrictions apply:

- You can have a single definition only for each channel group.
- You can have a single definition only for each slot-data port combination. This means that you cannot create an EtherChannel VLAN definition if a data port definition already exists for this slot-data port.

Related Topics

- Deleting EtherChannel VLAN Definitions, page 68-45
- Creating or Editing Data Port VLAN Definitions, page 68-46
- IDSM Settings, page 68-43

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > IDSM Settings** from the Policy selector.
- (Policy view) Select **Catalyst Platform > IDSM Settings**.

The IDSM Settings page is displayed. For a description of the fields on this page, see IDSM Settings Page, page 68-47.

Step 2 Do one of the following:

- To create an IDSM EtherChannel VLAN definition, click **Add Row** beneath the EtherChannel VLANs table.
- To edit an IDSM EtherChannel VLAN definition, select it in the list, then click **Edit Row** beneath the table.

The IDSM EtherChannel VLAN dialog box is displayed. For a description of the fields in this dialog box, see Create and Edit IDSM EtherChannel VLANs Dialog Boxes, page 68-49.

Step 3 To assign a channel group number to the Ethernet interface for the VLAN, or to change the channel group number, enter a number in the **Channel Group** text box.

Step 4 To associate the VLAN with the numbered chassis slot where you installed your IDSM services module and to associate one module data port with the VLAN, do one of the following:

- Enter the slot-port number in the **Slot-Ports** text box.
- Click **Select** to open the IDSM Slot-Port Selector dialog box.

Note Associating one module data port with the VLAN enables you to configure the port at the group level instead of configuring it manually.

Step 5 From the Mode list, select the running mode of the EtherChannel VLAN. If you select Capture, select the check box to configure the specified channel group as a capture destination.

Note If you do not select this check box, the capture port is created in shutdown mode.

Step 6 To include a VLAN in the specified channel group, do one of the following:

- Enter its numeric ID in the VLAN IDs text box.
- Click **Select** to open the VLAN Selector dialog box.

You can enter or select more than one VLAN ID.

Step 7 Click **OK** to save your definitions locally on the client and close the dialog box.

Deleting EtherChannel VLAN Definitions

You can delete an EtherChannel VLAN definition on the IDSM.

Related Topics

- Creating or Editing EtherChannel VLAN Definitions, page 68-44
- Deleting Data Port VLAN Definitions, page 68-47
- IDSM Settings, page 68-43

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > IDSM Settings** from the Policy selector.
- (Policy view) Select **Catalyst Platform > IDSM Settings**.

The IDSM Settings page is displayed. For a description of the fields on this page, see IDSM Settings Page, page 68-47.

Step 2 Click a row in the table to select the VLAN definition to delete.

Step 3 Click **Delete Row**.

Creating or Editing Data Port VLAN Definitions

When defining a data port VLAN definition, you must:

- Define the slot-port combination where the data port is located.
- Select the sensing mode used by the data port.
- Define which VLANs are forwarded to the data port.

The following restrictions apply:

- You may have a single definition only for each data port.
- You cannot create a data port definition if the port is already defined as part of a channel group.

Related Topics

- Deleting Data Port VLAN Definitions, page 68-47
- Creating or Editing EtherChannel VLAN Definitions, page 68-44
- IDSM Settings, page 68-43

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > IDSM Settings** from the Policy selector.
- (Device view) Select **Catalyst Platform > IDSM Settings**.

The IDSM Settings page is displayed. For a description of the fields on this page, see IDSM Settings Page, page 68-47.

Step 2 Do one of the following:

- To create an IDSM data port VLAN definition, click **Add Row** beneath the Data Port VLANs table.
- To edit an IDSM data port VLAN definition, select it in the list, then click **Edit Row** beneath the table.

The IDSM Data Port VLAN dialog box is displayed. For a description of the fields in this dialog box, see Create and Edit IDSM Data Port VLANs Dialog Boxes, page 68-49.

Step 3 To associate the VLAN with the numbered chassis slot where you installed your IDSM services module and to associate one module data port with the VLAN, do one of the following:

- Enter the slot-port number in the **Slot-Ports** text box.
- Click **Select** to open the IDSM Slot-Port Selector dialog box.

Note Associating one module data port with the VLAN enables you to configure the port at the group level instead of configuring it manually.

Step 4 From the Mode list, select the running mode of the data port VLAN. If you select Capture, select the check box to configure the specified data port as a capture destination.

Note If you do not select this check box, the capture port is created in shutdown mode.

Step 5 To assign a VLAN to the specified data port, do one of the following:

- Enter its numeric ID in the VLAN IDs text box.
- Click **Select** to open the VLAN Selector dialog box.

You can enter or select more than one VLAN ID.

Step 6 Click **OK** to save your definitions locally on the client and close the dialog box.

Deleting Data Port VLAN Definitions

You can delete a data port VLAN definition on the IDSM.

Related Topics

- Creating or Editing Data Port VLAN Definitions, page 68-46
- Deleting EtherChannel VLAN Definitions, page 68-45
- IDSM Settings, page 68-43

Step 1 Do one of the following:

- (Device view) Select a Catalyst device, then select **Platform > IDSM Settings** from the Policy selector.
- (Policy view) Select **Catalyst Platform > IDSM Settings**.

The IDSM Settings page is displayed. For a description of the fields on this page, see IDSM Settings Page, page 68-47.

Step 2 Click a row in the table to select the VLAN definition to delete.

Step 3 Click **Delete Row**.

IDSM Settings Page

Use the IDSM Settings page to view and configure the VLAN settings for data ports and channel groups on Intrusion Detection System Service Modules (IDSM).

Navigation Path

You can access this page from:

- (Device view) Select **Platform > IDSM Settings** from the Device Policy selector.
- (Policy view) Select **Catalyst Platform > IDSM Settings** from the Policy Types selector.

Related Topics

- Create and Edit IDSM EtherChannel VLANs Dialog Boxes, page 68-49

- Create and Edit IDSM Data Port VLANs Dialog Boxes, page 68-49
- Chapter 46, "Managing Firewall Devices"
- Filtering Tables, page 1-48

Field Reference

Table 68-22 IDSM Settings Page

Element	Description
EtherChannel VLANs table	
Channel Group	Identifies the EtherChannel group to which the Ethernet interface is assigned.
Module Slot-Data Port	Identifies the IDSM service module data port by number (1 or 2) to distinguish between the two ports. Each IDSM service module (blade) has two data ports. You can configure a data port individually or you can assign it to an EtherChannel group. All data ports in a channel group are configured at the group level
Mode	Indicates whether the running mode is trunk (IPS) or capture (IDS).
Capture Enabled	Indicates whether the specified channel group is configured as a capture destination.
Allowed VLANs	Lists which VLANs are allowed for the specified channel group.
Add Row button	Opens the Create IDSM EtherChannel VLANs dialog box. From here you can define which traffic is directed to the data ports in an EtherChannel group and which sensing mode is used.
Edit Row button	Opens the Edit IDSM EtherChannel VLANs dialog box. From here you can modify the attributes of an EtherChannel VLAN definition.
Delete Row button	Deletes the selected VLAN from the IDSM.
Data Port VLANs table	
Module Slot-Data Port	Identifies the IDSM service module data port by number (1 or 2), to distinguish between the two ports.
Mode	Indicates whether the running mode is trunk (IPS) or capture (IDS). To change the mode, select and edit the relevant table row.
Capture Enabled	Indicates whether the specified data port is configured as a capture destination.
Allowed VLANs	Lists which VLANs are allowed for the specified data port.
Add Row button	Opens the Create IDSM Data Port VLANs dialog box. From here you can define which traffic is directed to a specific data port and which sensing mode is used.
Edit Row button	Opens the Edit IDSM Data Port VLANs dialog box. From here you can modify the attributes of a data port VLAN definition.
Delete Row button	Deletes the selected VLAN from the IDSM.

Create and Edit IDSM EtherChannel VLANs Dialog Boxes

Use the Create IDSM EtherChannel VLANs dialog box (or the Edit IDSM EtherChannel VLANs dialog box) to configure or reconfigure the attributes of an IDSM EtherChannel VLAN.

Navigation Path

Go to the IDSM Settings Page, page 68-47, then click the **Add** or **Edit** button beneath the EtherChannel VLANs table.

Related Topics

- Create and Edit IDSM Data Port VLANs Dialog Boxes, page 68-49
- IDSM Slot-Port Selector Dialog Box, page 68-50
- Service Module Slot Selector Dialog Box, page 68-35

Field Reference

Table 68-23 Create and Edit IDSM EtherChannel VLANs Dialog Boxes

Element	Description
Channel Group	The EtherChannel group to which the Ethernet interface is assigned.
Slot-Ports (Select button)	Associates the chassis slot number (in which the relevant services module is installed) with the data port in the format *x* -*y*, where *x* is the slot number and *y* is the port number. For example, 2-1 refers to data port 1 in slot 2. Click **Select** to open the IDSM Slot-Port Selector Dialog Box, page 68-50. From here, you can select the IDSM slot-port combinations to include in the EtherChannel group.
Mode	The running mode of the EtherChannel group: • Capture (IDS)—The IDSM2 passively monitors network traffic that was copied to its data ports by the Catalyst switch using either VACL capture or SPAN. • Trunk (IPS)—The IDSM2 operates as an 802.1Q trunk by performing VLAN bridging between pairs of VLANs within the same data port.
Capture Enabled	Applies only when the running mode is Capture (IDS). When selected, configures the specified channel group as a capture destination. When deselected, the channel group does not act as a capture destination.
VLAN IDs (Select button)	Identifies which VLANs the specified channel group should allow. Click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select VLANs to include or exclude.

Create and Edit IDSM Data Port VLANs Dialog Boxes

Use the Create IDSM Data Port VLANs dialog box (or the Edit IDSM Data Port VLANs dialog box) to define which traffic is directed to an IDSM data port and which sensing mode is used on that traffic.

Navigation Path

Go to the IDSM Settings Page, page 68-47, then click the **Add** or **Edit** button beneath the Data Port VLANs table.

Related Topics

- Create and Edit IDSM EtherChannel VLANs Dialog Boxes, page 68-49
- IDSM Slot-Port Selector Dialog Box, page 68-50
- Service Module Slot Selector Dialog Box, page 68-35

Field Reference

Table 68-24 Create and Edit IDSM Data Port VLANs Dialog Boxes

Element	Description
Slot-Port	Associates the chassis slot number (in which the relevant services module is installed) with the data port in the format *x* -*y*, where *x* is the slot number and *y* is the port number. For example, 2-1 refers to data port 1 in slot 2. Click **Select** to open the IDSM Slot-Port Selector Dialog Box, page 68-50. From here, you can select the IDSM slot-port combinations to include in the data port VLAN definition.
Mode	The running mode of the data port: • Capture (IDS)—The IDSM2 passively monitors network traffic that was copied to its data ports by the Catalyst switch using either VACL capture or SPAN. • Trunk (IPS)—The IDSM2 operates as an 802.1Q trunk by performing VLAN bridging between pairs of VLANs within the same data port.
Capture Enabled	Applies only when the running mode is Capture (IDS). When selected, configures the specified channel group as a capture destination. When deselected, the channel group does not act as a capture destination.
VLAN IDs (Select button)	Identifies which VLANs the specified data port should allow. Click **Select** to open the VLAN Selector Dialog Box, page 68-35. From here, you can select VLANs to include or exclude.

IDSM Slot-Port Selector Dialog Box

Use the IDSM Slot-Port Selector dialog box to associate slot-port objects with EtherChannel groups.

Navigation Path

Go to the Create and Edit IDSM EtherChannel VLANs Dialog Boxes, page 68-49 or the Create and Edit IDSM Data Port VLANs Dialog Boxes, page 68-49, then click **Select** in the Slot-Port field.

Related Topics

- VLAN Selector Dialog Box, page 68-35
- Filtering Tables, page 1-48

Field Reference

Table 68-25 IDSM Slot-Port Selector Dialog Box

Element	Description
Available IDSM Slot-Ports list	Displays the available slot-port definitions.
Add >> button	Applies only when selecting slot-ports for EtherChannel VLANs. Adds IDSM slot-port objects that you selected in the Available IDSM Slot-Ports list to the Selected IDSM Slot-Ports list.
Remove << button	Applies only when selecting slot-ports for EtherChannel VLANs. Removes selected IDSM slot-port objects from the Selected IDSM Slot-Ports list.
Selected IDSM Slot-Ports list	Displays the IDSM slot-port objects that are selected for an association with a data port or an EtherChannel group.

PART 7

Monitoring, Reporting, and Diagnostics

Viewing Events

Event Viewer enables you to selectively monitor, view, and examine events from ASA (including ASA-SM), FWSM and IPS devices. Events are organized into views that you can filter or search to find events that interest you. You can create customized views and filters to fit your needs, or use the predefined views included in the application.

This chapter contains the following topics:

- Introduction to Event Viewer Capabilities, page 69-1
- Overview of Event Viewer, page 69-7
- Preparing for Event Management, page 69-27
- Managing the Event Manager Service, page 69-30
- Using Event Viewer, page 69-37
- Examples of Event Analysis, page 69-58

Introduction to Event Viewer Capabilities

Event Viewer monitors your network for syslog (system log) events from ASA and FWSM devices and security contexts and SDEE (Secure Device Event Exchange) events from IPS devices and virtual sensors. Event Viewer collects these events and provides an interface by which you can view them, group them, and examine their details.

Note Beginning with version 4.5, Security Manager enables you to forward syslogs to one local collector and two remote collectors. For more information, see Event Management Page, page 11-27.

Tip Event Viewer and its related applications, Report Manager and Health and Performance Monitor, are useful for operational monitoring and troubleshooting of certain types of Cisco devices in your network. These applications do not provide extensive event correlation, compliance reporting, long-term forensics, or the integrated monitoring of both Cisco and non-Cisco devices.

When working with IPS events, the Report Manager component of Cisco Security Manager reports events individually; the Event Viewer component of Cisco Security Manager displays alerts. In the Event Viewer component, the IPS Summarizer groups events into a single alert, thus decreasing the number of alerts that the IPS sensor sends out.

Tip Cisco IPS Manager Express (IME) and Cisco Security Manager do not summarize events in precisely the same way.

This section briefly describes some key activities that Event Viewer facilitates:

- Historical View, page 69-2
- Real-Time View, page 69-2
- Views and Filters, page 69-3
- Policy Navigation, page 69-3
- Understanding Event Viewer Access Control, page 69-4
- Scope and Limits of Event Viewer, page 69-4
- Deeply Parsed Syslogs, page 69-6

Historical View

An historical view is one that displays events from a selected period of time (for example, the last 10 minutes) and does not automatically update as new events are collected. You must refresh the view to see newer events.

Consider the following activities among the many possibilities for employing Event Viewer with an historical view:

- **Troubleshoot Connectivity**—When a report comes in that a user cannot reach a particular server, you can set an historical view (for example, the last 10 minutes) that displays all events that affect that user's IP address as a source or destination. Then, you can go from a particular displayed event to the policy denying that user's access to the resource.
- **Tune Signatures**—After setting a view of all IPS messages, or all IPS messages of a given category, you might decide that an event is actually a false positive. You can then cross launch into the associated policy and either tune the signature to exclude the host or lessen the reported severity of the particular event.

  Also consider creating an event action filter to modify how the alert is handled. Frequently, event action filters are a better way of dealing with false positives than editing the actual signature. For more information, see Configuring Event Action Filters, page 40-4.
- **Validate Policy Deployment**—After deploying a new or changed policy, you might want to confirm that it is operating effectively by selecting events corresponding to the given policy. For example, you could identify firewall-deny messages triggered by the new policy.

Real-Time View

A real-time view displays events as they are received and automatically updates the Event Table in waterfall fashion. Keep in mind that the term "real-time" is not precise. System latency and other factors prevent true real-time system response.

Consider the following activities among the many possibilities for employing Event Viewer with a real-time view:

- **Investigate Attacks in Near Real-time**—By isolating details of a particular source IP address, or a source/destination pair, Event Viewer can provide details about attacks on your monitored devices, or attacks that are going through those devices.
- **Validate Device Activity**—You can examine a device in your network and determine whether it is present and whether it is sending events.
- **View High Threat IPS Events**—You can filter a view to display all events that exceed a certain threat level. On a properly tuned IPS sensor, this should be a manageable flow of events to watch in a real-time view.

Views and Filters

When you view events in Event Viewer, you open a view. A *view* is a set of filters and other properties, including color rules, selected columns and their positions and widths, and the default time window, that let you define a subset of events. Views help to limit the scope of the events list so that you can more easily find what you are looking for.

Event Viewer includes a number of predefined views. Although you cannot change the filter rules for these views, you can create copies of the views and change the filter rules in your copy. Views you create are called custom views. For more information, see Creating Custom Views, page 69-41.

Using filters is key to getting the most from Event Viewer. You can distill from all the events being received a view of only the information that you need or want. You can use the various methods of filtering to reduce the events list, filtering lists that have already been filtered. The following list explains the general filtering features; for more information, see Filtering and Querying Events, page 69-43.

- **Time filters**—You can use time filters to limit the events that are loaded into your client as well as to limit the events displayed in the Event Table. With time filtering you can select predefined values, such as **the last hour**, or specify a particular time range by dates and times. For more information, see Selecting the Time Range for Events, page 69-43.
- **Column filters**—You can use column filters to filter events based on a particular value of an event. For example, you could filter on a particular source or destination, or both. For certain columns you can also filter on a range of values or on a policy object. Column filters are part of the view settings for a view. For more information, see Creating Column-Based Filters, page 69-45.
- **Quick filters**—You can use quick filters to execute a text-based filter on events listed in the event table. The search is not column-sensitive, showing all events in which the string appears in any column. You can use the Quick Filter drop-down list (shown as a magnifier) to modify the scope of the filter. For more information, see Filtering on a Text String, page 69-47.
- **Drilling down with filters**—Aggregating additional filters allows you to become more and more selective, to “drill down” until you can view a particular event or set of events that meet your requirements. The View Settings pane at the top of the Event Monitoring window updates with each additional filter choice you make to show the current aggregate filter definition of the view selected.

Policy Navigation

You can navigate from a particular event to the policy within Security Manager that governs that event. You can also navigate from certain policies to events associated with those policies. For more information, see Looking Up a Security Manager Policy from Event Viewer, page 69-54 and Looking Up Events for a Security Manager Policy, page 69-55.

Understanding Event Viewer Access Control

The user privileges assigned to your username control what you can do in Event Viewer. If you use local users, or other types of non-ACS access control, then all users have access to Event Viewer. However, the following access limits are imposed:

- You must have system administrator, network administrator, or approver privileges to select or deselect devices for monitoring. See Selecting Devices to Monitor, page 69-34.
- You must have system administrator privileges to change the Event Management administrative settings page, where you enable or disable the service and configure storage location and other settings, as described in Starting, Stopping, and Configuring the Event Manager Service, page 69-30 and Event Management Page, page 11-27

If you use ACS to control access to Security Manager, you can also control the following:

- You can control access to the Event Viewer application using the View Event Viewer privilege. Using this privilege, you could prevent certain users from accessing Event Viewer, or create roles that allow access to Event Viewer without allowing access to Report Manager. All default ACS roles are permitted to use Event Viewer.
- You can control which users can enable or disable monitoring for devices using the Modify > Manage Event Monitoring privilege. A user must have this privilege to select devices for monitoring as described in Selecting Devices to Monitor, page 69-34. The default ACS roles that have this permission are system administrator, network administrator, approver, security administrator, and security approver.
- You can control the use of the policy lookup feature. Users must have View Device privileges to the device, and also View privileges to the firewall or IPS policy, to perform policy lookup. If users do not have all permissions, they will get an “Unable to Find Matching Rule” error if they try to look up a matching rule. For more information about policy lookup, see Looking Up a Security Manager Policy from Event Viewer, page 69-54.
- Users can view events on devices only if they have at least View privileges to the device.
- You can control access to the Event Management administrative settings page, where you enable or disable the service and configure storage location and other settings, as described in Starting, Stopping, and Configuring the Event Manager Service, page 69-30 and Event Management Page, page 11-27. The user must have Admin privileges to access this page (or any other administrative settings page). All default ACS roles except help desk can view the page, but only system administrators can change settings.
- You can control the use of network/host and service policy objects for column filters (such as the Device, Source, Destination, Source Service, and Destination Service columns). Users must have the appropriate View Object permissions for network/host, network/host-IPv6, and service objects to use them in filters. For more information on creating column filters, see Creating Column-Based Filters, page 69-45.

For information on integrating Security Manager with Cisco Secure ACS, see the *Installation Guide for Cisco Security Manager*.

Scope and Limits of Event Viewer

The following table provides details on the functional scope and limits of Event Viewer:

Table 69-1 Event Viewer Scope and Limits

Item	Description
Device Support	You can view events collected from the following types of devices. Although Event Viewer has been tested with the indicated software releases, you might be able to use it with older software releases. • **ASA devices (including ASA-SM) and security contexts**—All 8.x releases. • **FWSM devices and security contexts**—Releases 3.1.17, 3.2.17, 4.0.10, and 4.1.1 and higher. • **IPS devices and virtual sensors**—Releases 6.1 and higher. IPS support does not include IOS IPS.
Event Data Store Size and Location	You can control the location and disk space allocated to holding events collected from monitored devices. After the Event Data Store is 90 percent filled, newest events replace oldest events. You can also configure an extended storage, or archive, location on attached storage devices. Security Manager automatically copies events into the extended storage; when you view historical events, they are automatically retrieved from extended storage if they no longer reside on the local disk. For more information on configuring these settings, see the Event Management Page, page 11-27
Event Limit	You can control the maximum number of events that can be viewed at one time in the events table using the Event Data Pagination Size option. For information on configuring the option, see Event Management Page, page 11-27.
Policy Objects	You can use some types of policy objects, such as network/host and services objects, when creating column filters. You can also view host object names instead of IP addresses in the source and destination columns by selecting **View > Show Network Host Objects**. This option is selected by default. IP address to host name mapping is supported only for the source and destination of events. Also, the mapping applies to Host objects only; Event Viewer will not show an object name when the source or destination of an event matches a Network object, Group object, or Address Range object. **Tip** Hover over a host object name to view the IP address associated with that object.
Views	A single Event Viewer client can open at most four historical views and one real-time view at a time.
Clients	For a single Security Manager server, a maximum of 5 Security Manager clients can open Event Viewer at one time, and a Security Manager client can open one copy of Event Viewer.

Deeply Parsed Syslogs

The structure and contents of standard syslogs and the elements comprised by each are detailed in the System Logs documentation for the device and software version you are using.

You can find the documentation on Cisco.com at these locations:

- ASA Devices: http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html
- FWSM Devices: http://www.cisco.com/en/US/products/hw/modules/ps2706/products_system_message_guides_list.html

Syslogs other than those listed here are presented as raw syslogs. Only deeply parsed syslogs present the full content carried by the syslog.

The deeply parsed syslogs in Security Manager are detailed in the following table.

Table 69-2 Deeply Parsed Syslogs

Syslog Category	Syslog ID	Total Number of Syslogs
Flow, Session Syslogs	110002-110003, 209003-209005, 302003-302004, 302009-302010, 302012-302018, 302020-302021, 302035-302036, 302303-302306, 302033-302034, 303002-302005, 313001, 313004, 313005, 313008, 324000-324006, 337001-337009, 431001-431002, 407001-407002, 416001, 418001-418002, 419001-419003, 424001-424002, 450001, 448001, 609001-609002 **Note** The 302303-302306 state-bypass syslog has been deep parsed only for event manager. However, the event description in the event manager for TCP, UDP, and SCTP state-bypass syslogs, does not display the "State-bypass" keyword. **Note** Reporting, Event to Policy, and Policy to Event are not supported for state-bypass syslog.	66
Botnet	338001-338004, 338101-338104, 338201-338202, 338301	11
ACL	106100, 106023, 106002, 106006, 106018	5
Denied Firewall	106001, 106007, 106008, 106010-106017, 106020-106022, 106025-106027	17
Identity Firewall	746003, 746005, 746010, 746016	4
AAA	109001-109010, 109012, 109016-109020, 109023-109029, 109031-109035, 113001-113025	53
Inspect	108002-108007, 303004-303005, 400000-400050, 406001-406002, 415001-415020, 500001-500005, 508001-508002, 608001-608005, 607001- 607003, 703001-703002, 726001	99
NAT	201002-201006, 201009-201013, 202005, 202011, 305005-305012	20

Table 69-2 Deeply Parsed Syslogs (continued)

Syslog Category	Syslog ID	Total Number of Syslogs
IPSec VPN	402114-402122, 602103-602104, 602303-602304, 702305, 702307	15
Failover (HA)	101001-101005, 102001, 103001-103007, 104001-104004, 311001-311004, 709001-709007, 210001-210022 (except 210008, 210010)	48
SSL VPN	725001-725009, 725012-725013, 716001-716020, 716023 -716039, 716041-716060, 722001-722023, 722026-722044, 722046-722051, 723001-723002, 723009-723012, 723014, 724001-724004	128
Etherchannel	426001-426003	3
Cluster	302022- 302027	6

Overview of Event Viewer

Use Event Viewer to view events and alerts collected from monitored firewall and IPS devices. For more information about selecting devices for monitoring, see Selecting Devices to Monitor, page 69-34.

To open Event Viewer, do any of the following:

- Select **Start > All Programs > Cisco Security Manager Client > Event Viewer** from the Windows Start menu (your exact command path might differ), or double-click the Event Viewer icon on the desktop. You are prompted to log in. For more information about starting a Security Manager client application, see Logging In to and Exiting the Security Manager Client, page 1-12.
- Select **Launch > Event Viewer** from the Configuration Manager or Report Manager applications, or click the Event Viewer button on the Configuration Manager toolbar. Event Viewer is opened using the same user account that you used to log into the other application.

The following illustration and subsequent list explain the basics of Event Viewer.

Figure 69-1 Event Viewer Main Window

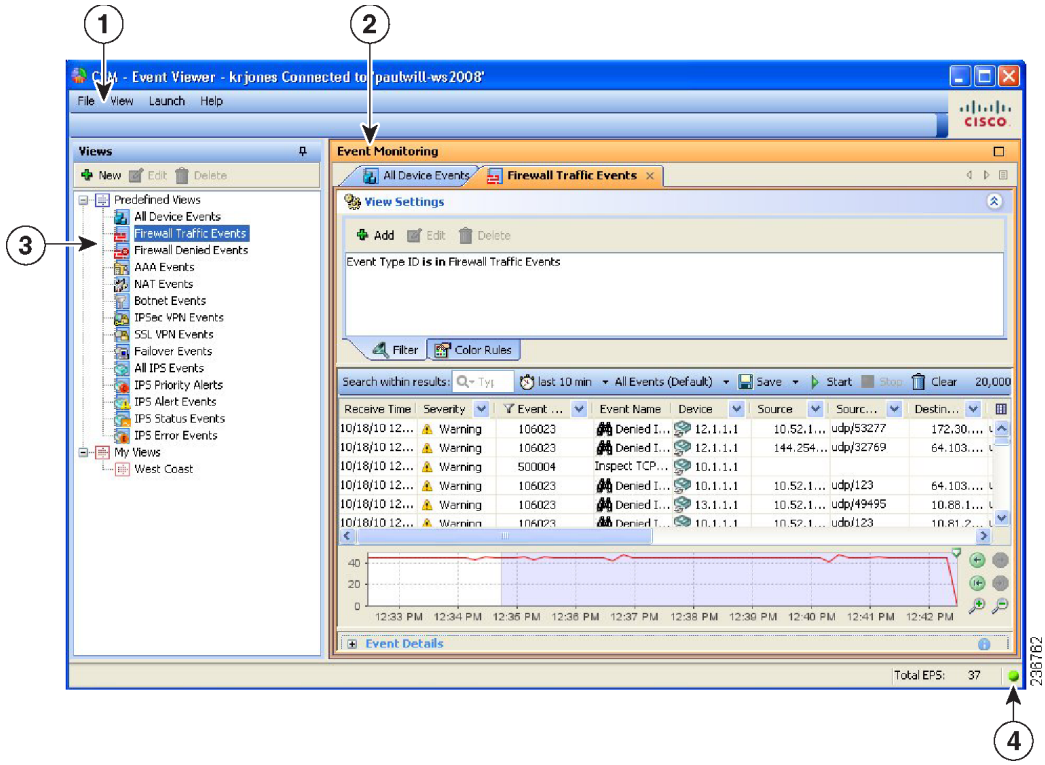

The following list explains the main Event Viewer window in more detail.

- **(1) Menu Bar**—General commands for performing actions in Event Viewer, including the following menus:
  - File, for operations on views. For information on the commands, see Event Viewer File Menu, page 69-9.
  - View, for operations within a view and general system management. For information on the commands, see Event Viewer View Menu, page 69-10.
  - Launch, for opening the Configuration Manager or Report Manager applications.
  - Help, for opening the online help or for viewing copyright and licensing information.
- **(2) Event Monitoring Window**—The right pane shows the open views. Each open view is represented on separate tabs (you can have up to four open historical views and one real-time view). Note that you can arrange views horizontally or vertically in this space, or even make a view float to a separate window. For more information about how you can arrange or float views, see Floating and Arranging Views, page 69-38.

  For detailed information about the many parts of the event monitoring window, see Event Monitoring Window, page 69-14.
- **(3) View List**—The left pane is a list of views. The list is organized into folders that separate the predefined views from the custom views, which are listed in the My Views folder. The simplest way to open a view is to double-click it, which replaces the currently open view. To open a view without replacing the currently open view, right-click the view and select **Open in New Tab**. For more information on opening views, see Opening Views, page 69-38.

  For information about other things you can do with the view list pane, see View List, page 69-12.

- **(4) Status Information**—The lower right portion of the status bar shows the current events per second (EPS) rate and an icon that indicates the current health of the monitoring system. Click the alert status icon to open a bubble that shows statistics for the past five minutes and any current system alerts. From this view, you can click the Details link to see more detailed information; click the alert status icon again to close the bubble. For more information, see Monitoring the Event Manager Service, page 69-31.

Note The Events per Second (EPS) information that is displayed on the Status Bar is calculated based on the number of events received every two seconds. Whereas, the EPS information that is displayed on the Time Slider graph is calculated by performing an aggregation of all the events that are available in a selected time range. Therefore, the numbers displayed on the Status Bar and the Time Slider graph might differ.

See an example below:

Example of EPS information display on the Status Bar

Assume at time T1, the Event Viewer application received 192 events. The Events per second (EPS) that will be displayed on the Status Bar is 192 / 2 = 96. This is because Security Manager collects events every 2 seconds and displays the Events per Second on the Status Bar. Let us say at T1 + 2 seconds, the Event Viewer application received 384 events. The EPS that will be displayed on the Status Bar is equal to (384 - 192)/2 = 96. This is the difference between current and previous value divided by 2.

Example of EPS information display on the Time Slider graph

Security Manager persists the events per second at an interval of 10 seconds. For example, if the Event Viewer application received 352 events in 10 seconds interval, the EPS is equal to 352 / 10 = 35. This value is persisted by Security Manager. For the next 10 seconds interval, if the Event Viewer receives 1056 events, the EPS will be equal to (1056 – 352)/ 10 = 70, which is persisted by Security Manager.

Displaying values on the Time Slider graph

The Time Slider graph displays the information for a period of time which has a Start Time and an End Time. All events per second that were collected in the given time interval are aggregated and plotted on the graph. In the given example, 35 and 70 are the values stored every 10 seconds. Therefore the Time Slider graph displays EPS as 35 and 70, which are different from the values displayed in the Status Bar.

Event Viewer File Menu

The following table describes the commands on the File menu in Event Viewer.

Table 69-3 File Menu in Event Viewer

Command	Description
New View	Creates a new custom view. You are prompted for a name and description. See Creating Custom Views, page 69-41. Alternatively, click the **New (+)** button in the view list.

Table 69-3 File Menu in Event Viewer (continued)

Command	Description
Open View	Opens a view on a new tab. You are prompted to select the view to open. You can open at most four historical views and one real-time view. See Opening Views, page 69-38. **Tip** You can double-click in the view list to open the view and replace the view that is displayed.
Save	Saves changes made to the active view, including filters (for custom views only), table preferences such as selected columns, column width, and sort order, the time range, and color rules. See Saving Views, page 69-42. If you want to save filter changes for a predefined view, you must use Save As to create a new custom view.
Save As	Saves as a custom view the changes you have made to the displayed view. See Saving Views, page 69-42.
Close View	Closes the displayed view.
Close All Views	Closes all open views.
Exit	Closes Event Viewer. Exiting the application closes any floating Event Viewer window that is open.

Event Viewer View Menu

The following table describes the commands on the View menu in Event Viewer.

Table 69-4 View Menu in Event Viewer

Command	Description
Mode	Specifies the time interval for selecting the events to display in the event table. From the submenu, select one of the following options: • **last 10 minutes** • **last 1 hour** • **last 12 hours** • **last 1 day** • **last 1 week** • **is today** • **is yesterday** • **is on . . .** (Opens a calendar on which you click to specify a single date.) • **is between** (Opens two calendars on which to specify a beginning and ending date and time.) • **Real Time** (Sets the mode to display events as they are received.) Alternatively, click the **Time Selector** control on the toolbar and select from the same options. See Event Table Toolbar, page 69-16.
Customize Column	Changes the columns shown in the event table. The Choose Columns to Display dialog box opens, where you can select the columns that you want to display. For details on the columns available, see Columns in Event Table, page 69-18.
Start	Initiates retrieving events to update the current view's event table. The event table then displays events received from the moment you clicked **Start** back to either the limit of the time mode or the event table pagination limit. Alternatively, click the **Start** button on the event table toolbar.
Stop	Stops event retrieval. The event table then displays the events received until the moment you clicked Stop. Alternatively, click the **Stop** button on the event table toolbar.
Show View Settings	Opens the View Settings pane, which displays the filters and color settings for the current view. You can alter these settings using the View Settings pane. Alternatively, click anywhere in the View Settings pane title bar, such as on the icon, on the text, or on the double arrow on the right side of the title bar. Clicking the heading opens and closes the pane.

Table 69-4 View Menu in Event Viewer (continued)

Command	Description
Show Event Details	Opens the Event Details pane and displays the selected event's details. Alternatively: • Click the expand icon (+) on the left of the Event Details pane title bar. • Double-click an event in the event table to display the event details data in a pop-up window. **Tip** From the Event Details dialog box you can print the event details or you can copy one or more of the detail rows to your clipboard. You can also scroll through the events list using the Next and Previous buttons.
Manage Monitored Devices	Allows you to select which devices, or groups of devices, can have events displayed in Event Viewer. For more information, see Selecting Devices to Monitor, page 69-34. **Note** By default, any ASA, FWSM, or IPS device added to the Security Manager inventory is monitored.
Show Event Store Disk Usage	Opens a window that displays the amount of disk space used as well as the age of the oldest event stored. See Monitoring Event Data Store Disk Space Usage, page 69-35.
Show Network Host Objects	When selected, the host object name is displayed, if available, instead of the Source or Destination IP address. This option is selected by default. **Tip** Hover over a host object name to view the IP address associated with that object.
Reset Layout	Re-establishes the width of the view list pane to its original setting after it has been hidden or manually narrowed or widened

View List

The left pane of the Event Viewer main window displays a list of available views as shown in the following illustration. A *view* is a set of filters and other properties, including color rules, selected columns and their positions and widths, and the default time window, that let you define a subset of events.

Figure 69-2 Event Viewer View List

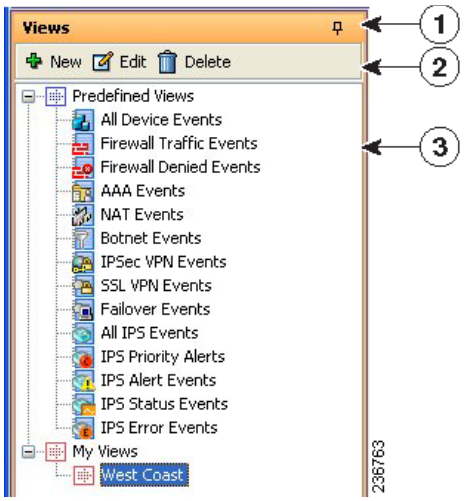

The view list includes the following controls:

- **(1) Push Pin button**—Click the Push Pin icon to control whether the view list pane is opened or closed. If the pin is vertical, the view list remains open unless you maximize the event monitoring window (the right pane). If the pin is horizontal, the view list collapses to the left margin, and you must click the Views heading in the left margin to open the list.
- **(2) Toolbar**—The toolbar contains these buttons:
  - **New button**—Click the New button to create a new custom view. You are prompted for a view name and description. For more information, see Creating Custom Views, page 69-41.
  - **Edit button**—Click the Edit button to change the name or description of the selected custom view. You can edit custom views only. For more information, see Editing a Custom View Name or Description, page 69-41.
  - **Delete button**—Click the Delete button to delete the selected custom view. You can delete custom views only. For more information, see Deleting Custom Views, page 69-43
- **(3) List of views**—The list is organized into folders that separate the predefined views from the custom views, which are listed in the My Views folder. The simplest way to open a view is to double-click it, which replaces the currently open view. To open a view without replacing the currently open view, right-click the view and select **Open in New Tab**. For more information on opening views, see Opening Views, page 69-38.
- **Right-click shortcut menu**—If you right-click on a view, you get a list of additional commands that you can perform:
  - Open—Opens the view and uses it to replace the currently active view. If the currently active view contains unsaved changes, you are prompted to save them. If the view is already open, it is brought to the foreground. See Opening Views, page 69-38.
  - Open in New Tab—Opens the view in a new tab, so that no existing open views are closed. See Opening Views, page 69-38.
  - Save As—Saves the view as a new custom view. See Saving Views, page 69-42.
  - Edit—Edits the custom view name and description. See Editing a Custom View Name or Description, page 69-41.
  - Delete—Deletes the custom view. See Deleting Custom Views, page 69-43.

- View Description—Displays the description for the view.

Related Topics

- Views and Filters, page 69-3
- Overview of Event Viewer, page 69-7
- Floating and Arranging Views, page 69-38

Event Monitoring Window

The Event Monitoring window shows the event views that you have opened. From this window, you can configure views and analyze and filter events.

Figure 69-3 Event Monitoring Window

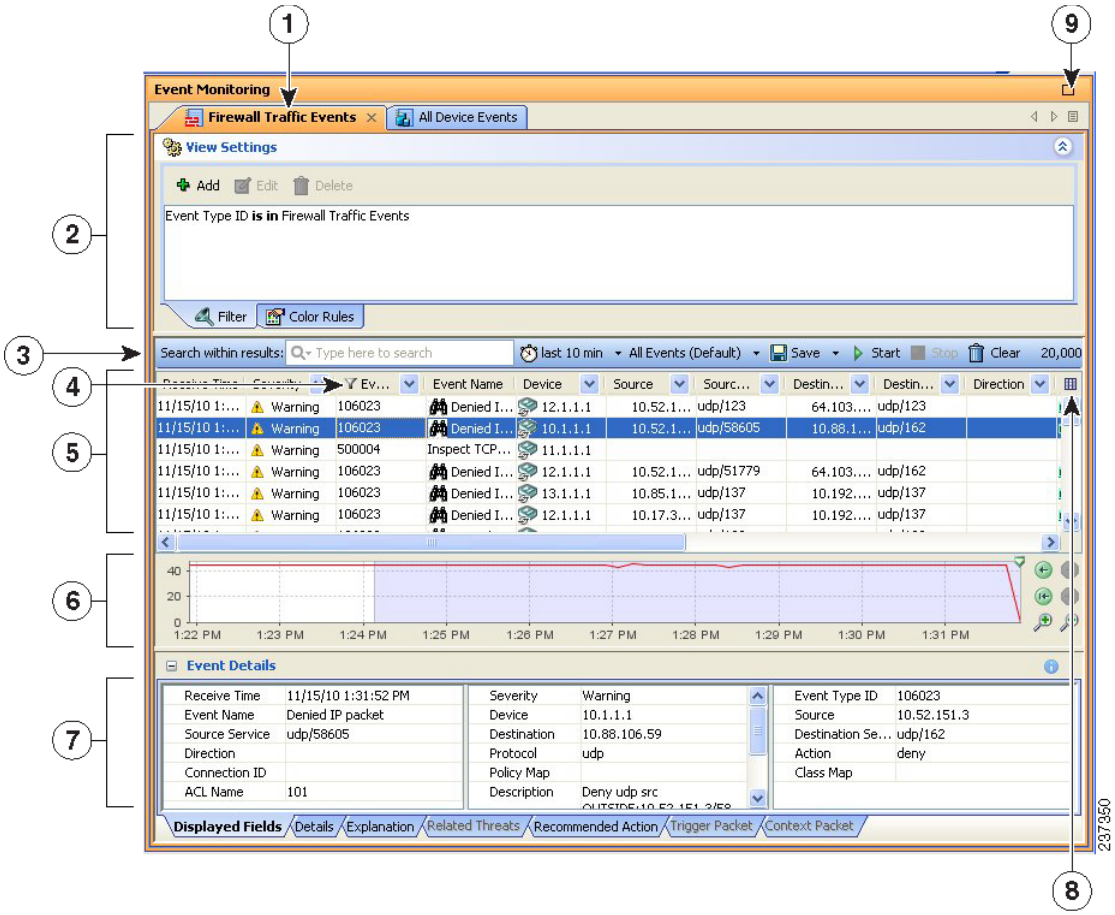

1	View tabs.	**6**	Time slider.
2	View Settings pane.	**7**	Event Details pane.
3	Event table toolbar.	**8**	Column selector button.
4	Filtered column icon.	**9**	Open view scroll buttons and list.
5	Event table.		

The Event Monitoring window contains these main elements:

- **View tabs (1, 9)**—When you open a view, it is represented as a tab in the window. To change views, click the tab, click the left or right arrow buttons to scroll through the tabs, or click the Open View List button and select the desired view. You can arrange views to be side by side or to float to separate windows by right-clicking the tab name and selecting an appropriate command; for more details, see Floating and Arranging Views, page 69-38.

Note You can open at most four historical views and one real-time view.

- **View Settings pane (2)**—Use the View Settings pane to define the column filters and color rules to use in a view. You can open and close the pane by clicking anywhere in the heading or by toggling the **View > Show View Settings** command.

  The View Settings pane contains two tabs: Filter and Color Rules. These tabs are shown along the bottom of the pane. On each tab, the body of the tab shows the current filter or rules; to change a rule, you select it and click the Edit or Delete buttons along the top of the pane, as appropriate. To create new rules, click the Add button.

  You can also add filters using the column filtering controls in the events table, as described in Creating Column-Based Filters, page 69-45. For more information on color rules, see Configuring Color Rules for a View, page 69-40.
- **Event Table Toolbar (3)**—The toolbar above the event table includes shortcut buttons and other controls that relate specifically to the events listed in the table. For a description of the toolbar controls, see Event Table Toolbar, page 69-16.
- **Event Table (4, 5, 8)**—The event table shows the events that match your filter criteria, one event per row. These events might be retrieved from the primary or the extended data store; you do not have to explicitly request data from the extended data store. To see events from a device, you must have View Device privileges to the device.

  The columns that make up the event table can be hidden, resized, reordered, and sorted upon as described in Customizing the Event Table Appearance, page 69-39. For a description of the columns, and how to use the column selector button to choose which columns are shown, see Columns in Event Table, page 69-18.

  If a column has a filter applied to it, an icon appears in the column heading.
- **Time Slider (6)**—For historical views, the time slider shows the current slice of time displayed in the table and the events per second rate as a linear graph. For more information about using the time slider, see Time Slider, page 69-25.
- **Event Details Pane (7)**—The event details pane shows detailed information about the currently selected event. You can open and close the pane by clicking anywhere in the heading or by toggling the **View > Show Event Details** command. For more information, see Event Details Pane, page 69-26.

Event Table Toolbar

The following illustration and table explain the elements in the toolbar that resides immediately above the event table in Event Viewer.

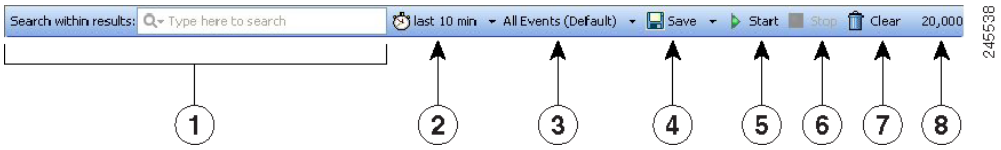

Table 69-5 Event Table Toolbar Elements

Callout	Name	Description
1	Search Within Results Field (Quick Filter)	This tool is also known as the *Quick Filter*. Use it to search for a word or phrase as well as to limit the scope of the search to certain columns. Further, you can select whether the search term used should be considered case sensitive, whether wildcards may be used, and whether a match may be partial, case sensitive, exact, or anywhere within a string. This search operates only on the selected view and within the data loaded. For more information, see Filtering on a Text String, page 69-47.
2	Time Selector (Mode) (Equivalent to **View > Mode**.)	You use the time selector to do the following: • Filter the events shown in the event table pane according to the time they were received. See Selecting the Time Range for Events, page 69-43. • Select between a real-time view or historical views. See Switching Between Real-Time and Historical Views, page 69-42. • Determine the time interval loaded in the client. If you are using one of the modes that shows events from the current time into the past, hovering the pointer over the field shows the start and end times for the displayed events. If you are using a specific time interval, the interval is shown in the toolbar.
3	Events by IP Address Type Selector	You use the events by IP address type selector to filter the list based on the type of IP address included in the events. Options are: • **All Events (Default)**—Show all events regardless of the address type. This is the default option. • **IPv4 Events Only**—Show events only if all addresses in the event are in IPv4 format. • **IPv6 Events Only**—Show events only if at least one address in the event is in IPv6 format. **Tip** You cannot save your selection. The next time you open the view, you need to reselect your option if you want something other than the default.
4	Save (Equivalent to **File > Save** or **File > Save As**.)	Click **Save** to save changes to the current view, including filters (for custom views only), table preferences such as selected columns, column width, and sort order, the time range, and color rules. Alternatively, click the down arrow and select **Save As** to save changes as a new custom view. If you want to save filter changes for a predefined view, you must use Save As to create a new custom view. For more information, see Creating Custom Views, page 69-41.

Table 69-5 Event Table Toolbar Elements (continued)

Callout	Name	Description
5	Start (Equivalent to **View > Start**.)	Click **Start** to reload or restart the listing of events in the Event Table. Clicking Start retrieves any events that have occurred since you originally loaded the table.
6	Stop (Equivalent to **View > Stop**.)	Click **Stop** to halt the listing of events in the Event Table. If you are in a real-time view, the Time Selector indicates the time stopped as well as the time interval that is loaded. Clicking on stop can also halt a query and display the set of events currently loaded in event viewer.
7	Clear	Click **Clear** to empty the event table.
8	Event Enumerator and messages	The number shown on the right of the toolbar indicates how many events are loaded onto the Event Viewer client. The number grows as events are loaded until either all events matching the filter criteria are displayed, or the pagination limit is reached, whichever is lowest. If you change the pagination limit (see Event Management Page, page 11-27), you must exit Event Viewer and open it again for the new limit to be in effect. If your query requires that events be retrieved from the extended event storage area, a message such as "Data being fetched from extended store" appears. Fetching events from the extended storage area typically takes longer than fetching them from the primary storage area.

Columns in Event Table

The following table lists alphabetically, and describes, all the columns that you can display in a view in Event Viewer. The columns applicable to a particular device vary, as does the presence or absence of event data for a particular event type.

When you save a view, the columns you selected, and their order, are preserved and displayed the next time you open the view. To select which columns to display in the open (and active) view, do one of the following:

- (Preferred method.) Click the **Column Chooser** icon in the far right of the event table header row (see Event Monitoring Window, page 69-14). The Choose Columns to Display dialog box that opens lists the columns in alphabetical order. You can select or deselect the columns either individually or by using the **Select/Unselect All** check box. Also, you can click Revert to return to the default column selection for the view.
- Select **View > Customize Columns**.
- Right-click any column heading and select or deselect a column individually, or select More to open the same dialog box that is used by the View > Customize Columns command.

Note Most columns other than Description, Event Name, and Receive Time include a filtering function. For more information, see Creating Column-Based Filters, page 69-45.

Table 69-6 Event Viewer Column Descriptions

Column Label	Description
AAA Group	The AAA group policy.
AAA Server	The server that handles user requests for access; it performs authentication, authorization, and accounting.
AAA User	The AAA username.
ACE Hash1 ACE Hash2	The hashcode1 and hashcode2 of the access control list entry (ACE). Hash codes are required for successful policy lookups from syslog 106023 and 106100 events. These hash codes are available only if you deployed the configuration using Security Manager.
ACL Name	The name or ID of the access control list (ACL).
Action	The action performed on the flow. For example: Terminated or denied.
Alert Details	The details regarding the alerts.
App Name	The name of the application originating the event.
App Stop Reason	An explanation of how or why the application was shut down.
App Version	The version of the application originating the event.
Attack Relevance Rating	A numerical value used to indicate an attack's relevance to its destination target or targets.
Backplane Interface	The backplane interface, which is identified only when the backplane interface differs from the physical interface.
Botnet Category	The category showing the reason a domain name is blocked, for example, botnet, Trojan, spyware, and so on.
Botnet Domain	The domain name or IP address in the dynamic filter database to which the traffic was initiated. It can be blocked, allowed, or grey listed.
Build Time	The date and time of the software build.
Build Type	The type of build. Typically this is a word such as "release" or "debug." In some cases, it is the ID of the builder of the application.
Byte Count	The number of bytes in the data transfer of the connection.
Call Id	The peer's Call ID for the session to which this packet belongs.
Class Map	The class map name.
Connection Duration	The lifetime of the connection.
Connection ID	A unique identifier for the connection.
Connection Limit	The maximum number of connections or sessions.
Connection Termination Value	A factor for which the connection is terminated, for example, incorrect version or invalid payload-type.
Current Connection Count	The number of current connections.
Description	For syslogs this shows the raw message, for IPS it shows a description of the event.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
Destination	The IP Address or hostname of the traffic destination (for ASA and FWSM) or the attack target (for IPS). It can be multi-valued and contain IPv4 or IPv6 addresses. If View > Show Network Host Objects is selected and a host object is defined that matches the destination IP address, the host object name is displayed. **Tip** Hover over a host object name to view the IP address associated with that object.
Destination Context Data	Context buffer indicating the data that was sent just prior to, and immediately after, the alert was triggered. A Base64-encoded representation of the stream data that was sourced by the target.
Destination FQDN	The fully-qualified domain name of the destination IP address, if any.
Destination Interface	The destination interface. For Etherchannel alerts (426001-426003), this is the name of the Etherchannel interface for which this event occurred. The member interface is identified in the Source Interface column.
Destination Locality	Whether the target address is located inside or outside of a given network as specified by the intrusion.
Destination OS	The target's operating system information.
Destination OS Relevance	A numerical value indicating the relevance of the destination target OS value.
Destination OS Source	The source of the Target OS data. Possible values are: learned, imported, or configured.
Destination Service	The destination port. It can be multi-valued.
Destination User Identity	The user name for the traffic destination, if any.
Device	The source of the event; usually the device ID. A device identified as Not Available has been deleted from the Security Manager inventory.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
Device Identifier	For a cluster of ASA devices, the ID of the event's source node, which is based on the "Enable Syslog Device ID" configuration on the Server Setup Page, page 54-21. You can use a "Device Identifier" to filter the syslogs generated by a failover device. In the event of a failover, the IP address of the failover device, that generated the syslog messages is displayed here. However, the Device Identifier column will be blank for syslog messages generated by failover device managed in Cisco Security Manager. **Note** Enable the **Process Syslogs from Failover Standby Device** check box in the **Event Management Page** in **Tools > Cisco Security Manager Administration**. A cluster is managed by Security Manager as a single device with multiple nodes. Thus, all the node's events are mapped to the cluster virtual IP and are displayed with the cluster virtual IP in Event Viewer. You can use "Device Identifier" to filter the syslogs generated by a specific cluster member of a node.
Direction	The direction of the traffic: inbound or outbound.
Event ID	A unique sequential number for each event, assigned internally.
Event Name	A user-friendly name given to the event.
Event Summary	Specifies that this is a summary alert, representing one or more alerts with common characteristics. The numeric value indicates the number of times the signature fired since the last summary alert with a matching initialAlert attribute value.
Event Type ID	For ASA or FWSM, the syslog ID. For IPS, this value could be: • A combination of Sig Id & Sub-Sig ID (for IPS Alert Events) • IPS Status (for IPS Status Events) • IPS Error (for IPS Error Events).
Execution State	The execution status of the application.
Final Alert	Applies to a summary alert, representing one or more alerts with common characteristics. It indicates whether this is the last event alert containing the same value in the initialAlert attribute.
Generation Time	Represents device local event generation time (available only for IPS events).
Global Correlation Audit Mode	Whether the alert was handled with audit mode processing: true or false.
Global Correlation Deny Attacker	Whether a deny-attacker action occurred (or would have occurred) because an internal override was exceeded due to the calculated risk rating: true or false.
Global Correlation Deny Packet	Whether a deny-packet action occurred (or would have occurred) because an internal override was exceeded due to the calculated risk rating: true or false.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
Global Correlation Modified Risk Rating	Whether the risk rating was adjusted by adding the reputation risk delta due to the risk rating: true or false.
Global Correlation Other Overrides	Whether any other defensive actions were taken because an override threshold was exceeded due to the calculated risk rating: true or false.
Global Correlation Risk Delta	A value from 0 to 99 that indicates how much the risk rating was increased due to the reputation score. If audit-mode is enabled, then it indicates how much the risk rating would have been adjusted had audit-mode not been enabled.
Hit Count	The number of times the flow was permitted or denied by the ACL entry in the configured time interval. The value is 1 when the ASA or FWSM generates the first syslog message for a particular flow.
Hit Count Info	ACL Hit Count information, for example, *First hit*.
Host ID	The globally unique identifier for the host that originated the event.
ICMP Code	The code of the ICMP type. For example, ICMP Type 3 and Code 0 is Net Unreachable or Code 1 is Host Unreachable.
ICMP Type	The type of ICMP message. For example, 3 for Destination unreachable, 8 for Echo.
Initial Alert	This field applies to a summary alert, representing one or more alerts with common characteristics. The value of InitialAlert provides the event ID of the last non-summary evIdsAlert with the same characteristic (sigid/subsigid).
Ip Log ID	The IP Log Identifier that uniquely identifies (with host-scope) an iplog document.
IpLog Address	The IPv4 or IPv6 address associated with the IP log.
IpLog Alert Reference	The global event ID of the evAlert event that triggered the log to be initiated.
IpLog Begin Time	The start of the time range that is currently available in the log document.
IpLog Bytes Captured	The total bytes captured. Note that some packets that were captured may have already been deleted from the log due to memory limitations.
IpLog Bytes Remaining	The number of bytes remaining until the log will be terminated.
IpLog End Time	The end of the time range that is currently available in the log document.
IpLog Minutes Remaining	The minutes remaining until the log will be terminated.
IpLog Packets Captured	The total number of packets captured and logged.
IpLog Packets Remaining	The packets remaining until the log will be terminated.
IpLog Status	A string that represents the log status.
IPS Category	The SEE event category.
IPS User	The username of the user initiating the operation.
License Limit	The maximum number of licenses.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
List Name	The list that includes the domain name, administrator allow list, block list, or IronPort list.
Login Action	The login action that occurred: loggedIn, loggedOut, or loginFailed.
Malicious Host	The hostname of the malicious host.
Malicious IP	The IP address of malicious device.
Max Connection	The maximum number of NAT connections.
MaxEmbryonic Connection	The maximum number of embryonic connections.
NAT Destination	The translated (also called natted) destination IP address. The host name of the translated destination.
NAT Destination Service	The translated (or natted) destination port.
NAT Global IP	The global address. It can contain IPv4 or IPv6 addresses.
NAT Source	The translated (or natted) source IP address. It can contain IPv4 or IPv6 addresses. The host name of the translated source.
NAT Source Service	The translated (or natted) source port.
NAT Type	The type of network address translation, for example *Static* or *Dynamic*.
New Time	The time to which the device clock was changed.
New Version	The system software version after an upgrade installation.
No.	The number of the event (row) in the current display. This is a simple sequential number and is not related to the content of the event. See the Event ID and Event Name fields for information about the type of event.
Old Time	The device clock time prior to a change.
Old Version	The system software version before an upgrade was uninstalled.
Operation Successful	Indicates whether an operation was successfully performed.
Package File	The name of package file to be auto-downloaded and installed.
Physical Interface	The physical interface, which is identified only if physical interface is different from the respective value in the Interface column.
Policy Map	The policy map name.
Protocol	The Level-3 or Level-4 protocol.
Protocol Version	The protocol version.
Protocol (Non L3)	Some non-Level-3 or -4 protocol seen in the event, for example, TACACS, RADIUS, FTP, or H245.
Reason	A rationale associated with certain events. For example, a connection tear down may have an associated reason.
Receive Time	The time the event was received by Security Manager.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
Reputation	The attacker's reputation score in the range -10.0 to +10.0. A lower (more negative) score indicates a greater likelihood that the host is malicious.
Result Status	The status of an operation, which indicates whether the operation successfully completed.
Risk Rating	A value that represents the calculated risk associated with the event.
Role in Cluster	The role of this member of an ASA load-balancing cluster: Cluster, Control Unit, or Data Unit.
Security Context	The security context with which the named interface, specified in the corresponding Interface column, is associated.
Sensor Event ID	The serial number for an event, which is guaranteed unique within the scope of the originating host.
Severity	The firewall or IPS severity values.
SIA Event Name	The event that occurred for the service identified in the SIA Service Name field.
SIA Service Name	The name of the Service Insertion Architecture (SIA) service for which this event occurred.
Sig Details	The details of the reported signature that was triggered and resulted in the generation of the alert.
Sig ID	The Sig ID value is used by the alert originator to identify the activity. It identifies the pre-defined signature defined for this activity.
Signature Version	The version of the signature definition used to generate an alert.
Source	The IP Address or hostname of the traffic source (for ASA and FWSM) or the attacker (for IPS). It can be multi-valued and contain IPv4 or IPv6 addresses. If View > Show Network Host Objects is selected and a host object is defined that matches the source IP address, the host object name is displayed. **Tip** Hover over a host object name to view the IP address associated with that object.
Source Context Data	The context buffer indicating the data that was sent just prior to and immediately after the alert was triggered. A Base64-encoded representation of the data stream that was sourced by the attacker.
Source FQDN	The fully-qualified domain name of the source IP address, if any.
Source Interface	The source interface. For Etherchannel alerts (426001-426003), this is the name of the interface that is part of the Etherchannel bundle for which this event occurred. The Etherchannel interface is identified in the Destination Interface column.
Source Locality	Identifies whether the attacker address is located inside or outside of a given network, as specified by the intrusion detection device's configuration.

Table 69-6 Event Viewer Column Descriptions (continued)

Column Label	Description
Source Service	The source port.
Source User Identity	The username associated with the traffic source, if any.
SSO Server	The single sign-on (SSO) server name.
SSO Server Type	The single sign-on (SSO) server type, for example, SiteMinder.
Sub SigId	The sub-sig ID value, which is used by the alert originator in combination with the signature ID (sigId) to identify the activity.
Summary Type	Defines the common characteristics of all alerts in a summary alert.
Target Value Rating	The asset values associated with targets identified in alerts.
Threat Level	Shows one of the following values, if any threat level pertains: none, very-low, low, moderate, high, or very-high.
Threat Rating	The threat rating of the event, if any.
Time Zone	The local time zone at the originating host's location.
Translated Call ID	The peer's Translated Call ID for the session to which this packet belongs.
Trigger Packet	The single, complete packet (in base64 binary format) that triggered the alert.
Truncated	Whether the trigger packet contained in the event is truncated.
Tunnel Type	The VPN tunnel type.
Type	The AAA type, for example authentication, authorization, or accounting.
Upgrade Name	The name of the upgrade package that was uninstalled.
URI	The URI of the auto-upgrade server directory.
UTC Offset	The offset attribute of sensor local time indicates the number of minutes that must be added to the UTC time to convert to local time at the originating host.
Virtual Sensor	The name of the virtual sensor associated with the event.
VLAN Id	The VLAN number associated with packets involved in the activity that triggered the alert.
VPN Group	The VPN group policy.
VPN IPSec SPI	The IPSec Security Parameter Index.
VPN User	The VPN username.
Watchlist Delta	The amount that the risk rating value was increase due to the source of the activity associated with the alert being on a watchlist.

Time Slider

The time slider resides below the events table when you are using an historical view; it is not used with real-time views. The following illustration explains the time slider; the pagination controls on the right are explained in Table 69-7.

Figure 69-4 Time Slider Elements

You can use the time slider to do the following:

- View the EPS (events per second) trends for the server. You can use the controls on the right to zoom in or out to get a better view of the EPS trends over the time frame that concerns you.

  You can also click and drag the background of the time slider to position the time ranges within the window. Moving the background does not affect the selected time range.

- Select a slice of time for displaying events in the table, either by moving the vertical slider or by using the pagination controls. The position of the vertical slider determines the most recent event displayed in the event table. Whenever you alter the time slice, the event table is reloaded with the events that match the time period.

  The time range of the events displayed in the event table is determined by the selected time interval. For more information, see Selecting the Time Range for Events, page 69-43.

The following table explains the pagination controls to the right of the time slider.

Table 69-7 Time Slider Paging Controls

Element	Description
	Previous page (earlier) and next page (later). The size of page varies according to the selected time mode. **Note** Using the page controls alternately, for example forward and then back, causes the sort order in the event table to reverse. That is, the latest events go from the top of the table to the bottom, or from the bottom to the top.
	First page (earliest) and last page (most recent).
	Zoom in (smaller total time interval shown) and zoom out (greater time interval shown). Zooming does not change the content of the event table. A shaded blue area indicates the time interval currently displayed in the event table.

Event Details Pane

The Event Details pane (illustrated in Event Monitoring Window, page 69-14) presents information contained within a single event. The information, which is displayed in multiple tabs within the pane, varies according to the richness of the event and the capability of Event Viewer to parse the data. Components include:

- **Displayed Fields Tab**—Displays the fields shown in the Event Table.
- **Details Tab**—Displays all available fields for the selected event. The fields are presented in alphabetical order.

- **Explanation Tab**—Displays a generic explanation for this event type.
- **Related Threats Tab**—Displays threats correlated to the event. (IPS Events only.)
- **Recommended Action Tab**—Displays a generic recommendation for an event of this type. (Syslogs only.)
- **Trigger Packet Tab**—Displays trigger packet data. (IPS Events only.)
- **Context Packet Tab**—Displays context packet data from the Source (Attacker) and Destination (Target). (IPS events only.)
- **Notes**—Enables you to add a note so that you can revisit particular signatures later to see what you or other users have added for a signature or an event. For more detailed information, refer to Signatures Page, page 39-4

Note The notes added here will be persistent when you cross-launch to Configuration Manager.

Note If you have more than one event for a particular signature, then annotating one event will annotate all the events related to that signature.

Preparing for Event Management

Before you can view events generated from a device, you must configure the device to work with Event Viewer.

Ensuring Time Synchronization

Standard network management practice includes consideration of time differences and network device synchronization. Typically, this includes the use of a Network Time Protocol (NTP) server. Event Viewer is most easily used with a common time standard. However, it is worth noting that you can view the time an event is received by Security Manager (Receive Time), and for IPS devices, the time the event was generated by a device (Generation Time).

Whenever possible, configure the Security Manager server and the devices it is monitoring with the same NTP server.

Differences between the clock on the Security Manager server and the clock on a Security Manager client at the time the client is opened are taken into account while translating/mapping event data from the server time to the client time. If the time difference changes dynamically due to the Security Manager server time moving ahead, for example, the data retrieved from the server will show an updated timestamp, but the client will continue to map the time difference based on the times on the server and the client when the client was opened. In such a situation, no data will be seen on Event Viewer for a brief interval that corresponds to the time change on the server. For this reason, we recommend that Security Manager server clock changes are done less frequently and at a time that would be the least impacting.

Configuring ASA and FWSM Devices for Event Management

Note From version 4.17, though Cisco Security Manager continues to support FWSM features/functionality, it does not support any bug fixes or enhancements.

Before you can use Event Viewer, or any other application that analyzes syslog events, to view events generated from an ASA (including ASA-SM) or FWSM device, you must configure the logging policies on the device to generate and transmit syslog messages.

Note A cluster device with a Virtual IP address (beginning with Security Manager version 4.4) can be added if configured in both Security Manager and on the virtual device.

Tip Although you can configure devices individually to specify the appropriate logging configuration, it is likely that more than one ASA or FWSM device in your network would use the same logging configuration. Although this topic describes how to configure an individual device, you can also create shared policies and assign them to multiple devices. For more information about configuring and assigning shared policies, see Creating a New Shared Policy, page 5-54 and Modifying Policy Assignments in Policy View, page 5-54.

Besides the logging configuration described here, you can also configure logging for individual access control entries when you configure them either in firewall policies or ACL policy objects. The default is to log denied access only, but you can configuring ACL logging options to provide increased logging.

Note To reliably report events from contexts in multiple-context mode, Cisco Event Viewer requires an IP address for the management interface of each context.

Step 1 (Device view) Select the ASA or FWSM device or security context, then select **Platform > Logging > Syslog > Logging Setup** from the Policies selector.

In the policy, select **Enable Logging**. You can configure other options as needed. For detailed information about the options, see Logging Setup Page, page 54-15.

Step 2 Select **Platform > Logging > Syslog > Syslog Servers**.

Add the Security Manager server's IP address to the syslog servers table. Configure the server to use the UDP protocol. The default port, 514, is correct unless you configure a different port on the Security Manager Administration Event Management Page, page 11-27.

If you are using other event management applications, such as CS-MARS, also add those servers to this policy.

Note You can use EMBLEM message format if you desire; both traditional and EMBLEM formats are supported. Keep in mind that EMBLEM is not supported by CS-MARS, so do not send EMBLEM-formatted messages to a CS-MARS server.

For detailed information about the options in the Syslog Servers policy, see Syslog Servers Page, page 54-27.

Step 3 If you want to configure non-default syslog server settings, such as adding time stamps to syslog messages, changing the severity level of messages, or suppressing the generation of specific messages altogether, configure the **Platform > Logging > Syslog > Server Setup** policy. For detailed information, see Server Setup Page, page 54-21.

Step 4 (Optional)You can configure the **Platform > Logging > Syslog > Logging Filters** policy to fine-tune the kinds of messages sent to syslog servers. For detailed information about this policy, see Logging Filters Page, page 54-12 and Edit Logging Filters Dialog Box, page 54-13.

Following are some tips for configuring this policy:

- When adding the logging filter, select **Syslog Servers** for Logging Destination.
- You can create a simple filter based on message severity, or you can configure a much more complex filter based on event classes. If you elect to use event classes, you can do the configuration directly in the Logging Filters policy, or you can configure event lists separately in the **Event Lists** policy (see Event Lists Page, page 54-9).

Step 5 (Optional) You can configure the **Platform > Logging > Syslog > Rate Limit** policy to limit the quantity of messages generated per time interval, either by message severity or message number. This can help you avoid flooding the syslog server. See Rate Limit Page, page 54-18.

Step 6 (Optional, but recommended) You can configure the **Platform > Device Admin > Server Access > NTP** policy to specify a network time protocol server for ASA devices. Using NTP ensures consistent date and time information for easy event correlation. Specify the same NTP server you use for the Security Manager server. If you use different servers, ensure the servers are synchronized. See NTP Page, page 52-21.

Configuring IPS Devices for Event Management

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Before you can use Event Viewer to view events generated from an IPS device, you must configure the Allowed Hosts policy on the device to allow the Security Manager server access to the device. Because Security Manager also needs to be configured in the Allowed Hosts policy to allow configuration access, your IPS devices might already be configured correctly. You should also configure the network time protocol (NTP).

Configure the following policies for IPS devices in Device view to enable effective event management on those devices:

- **Platform > Device Admin > Device Access > Allowed Hosts**—(Required) Add the Security Manager server to the table. You can either identify the Security Manager server by its host IP address (for example, 10.100.10.10), or you can specify the network that it is on (for example, 10.100.10.0/24).

  If you are using other event management applications with the device, such as CS-MARS, ensure that you also add those servers to the policy.

  For more information about configuring the Allowed Hosts policy, see Identifying Allowed Hosts, page 36-7.

- **Platform > Device Admin > Server Access > NTP**—(Recommended) Configure the same NTP server that you use for the Security Manager server to ensure consistent date and time information for easy event correlation. If you use different servers, ensure the servers are synchronized. For more information, see Identifying an NTP Server, page 36-23.

Tip Although you can configure devices individually to specify the appropriate allowed hosts and NTP configuration, it is likely that more than one IPS device in your network would use the same configuration. Although this topic describes how to configure an individual device, you can also create shared versions of these policies and assign them to multiple devices. For more information about configuring and assigning shared policies, see Creating a New Shared Policy, page 5-54 and Modifying Policy Assignments in Policy View, page 5-54.

Managing the Event Manager Service

The Event Manager service enables the use of the Event Viewer application. For Event Viewer to function, the service must be started. There are several tasks that you can perform to configure and manage the overall functioning of the service.

This section contains the following topics:

- Starting, Stopping, and Configuring the Event Manager Service, page 69-30
- Monitoring the Event Manager Service, page 69-31
- Selecting Devices to Monitor, page 69-34
- Monitoring Event Data Store Disk Space Usage, page 69-35
- Archiving or Backing Up and Restoring the Event Data Store, page 69-36

Starting, Stopping, and Configuring the Event Manager Service

The Event Manager service must be running for you to use Event Viewer or Report Manager.

When you install Security Manager, the Event Manager service is automatically enabled unless the server does not meet the minimum memory requirements that are documented in the *Installation Guide for Cisco Security Manager*. Although you can manually start the service on a system that does not meet the minimum memory requirements, you might find the performance to be dissatisfactory. The key factors are the number of devices managed and their rate of event generation.

Tip If you get a message that Event Viewer is unavailable when you select Launch > Event Viewer, but the **Enable Event Management** option is selected in the Tools > Security Manager Administration > Event Management page, try restarting the Event Manager Service. First, deselect the Enable option and click Save. Wait for the service to stop. Then, select the Enable option, click Save, and wait for the service to finish restarting. You can then try opening Event Viewer again.

The following procedure explains how to start, stop, and configure the Event Manager service.

Related Topics

- Monitoring Event Data Store Disk Space Usage, page 69-35

Step 1 In the main Security Manager window (not Event Viewer), select **Tools > Security Manager Administration** and select **Event Management** from the table of contents.

Step 2 Do one of the following:

- To enable, or start, the Event Manager service, select **Enable Event Management**.
- To disable, or stop, the Event Manager service, deselect **Enable Event Management**.

You can also change the other settings, such as the event data store location and maximum size, the syslog port to which devices should send events, and the pagination size (which determines the maximum number of events loaded into the event table). You can also configure an extended event storage location to augment your primary storage location. For detailed information about these settings, see Event Management Page, page 11-27.

Note Beginning with version 4.5, Security Manager enables you to forward syslogs to one local collector and two remote collectors. For more information, see Event Management Page, page 11-27.

Step 3 Click **Save** to save your changes.

If you changed the Enable Event Management option, you are prompted to confirm that you want to start or stop the Event Manager Service. If you click **Yes**, the service is started or stopped immediately, and you are shown a progress indicator and told when the change is completed. Wait until the status change is completed before continuing.

If you change other settings, with the exception of the pagination size, the Event Manager service must be briefly stopped and restarted. You are shown a progress indicator.

Monitoring the Event Manager Service

The Event Manager service processes incoming syslog messages and retrieves SDEE alerts from monitored IPS devices. The amount of data processed varies depending on network activity. There can be times when the events per second (EPS) generated in the network is higher than the service can handle, in which case the service goes into throttle mode, selectively dropping events.

You can monitor the status of the service to identify congestion and address problems that arise. The status of the service is shown in an icon in the lower right corner of the status bar in Event Viewer, as shown in Overview of Event Viewer, page 69-7. The Total EPS indicates the current events per second rate that the service is experiencing. The alert status icon color indicates the following:

- Green dot—There are no problems. All events are being processed normally.
- Yellow dot—There are some warnings, for example, low severity events are being dropped.
- Orange dot—There are more serious issues, for example, low and medium severity events are being dropped.
- Red dot—There is a critical situation, for example, high severity events are being dropped or there is a significant problem with the system, such as problems with the syslog port or with the event data store location.

- Disconnected network wire—The Event Manager service is disabled, either intentionally or due to some server problem; no events are being stored or retrieved. If this is not intentional, restart the Event Manager service as described in Starting, Stopping, and Configuring the Event Manager Service, page 69-30.

To view detailed information, click on the alert status icon. A bubble opens that shows summary statistics for the past five minutes, including the number of events received and dropped and event server alert messages, if any. Click the alert status icon again to close the bubble.

When the bubble is open, you can click the **Details** link in the bubble to view more detailed information. Clicking the Details link opens the Event Statistics Details dialog box, which shows the following information:

- **Last 5 Minutes Statistics:**
  - **Events Received**—The total number of syslog events Received and SDEE alerts retrieved in the past five minutes by the service.
  - **Events Dropped**—The total number of events or alerts that the service had to drop due to congestion. This number indicates drops from monitored devices only, so the number should be zero in normal circumstances. A non-zero number indicates that the service is in throttle mode; look for messages in the Event Server Alerts section.
  - **Events from Unmonitored Devices**—The number of syslog messages sent to the server that came from devices that are not selected for monitoring (as described in Selecting Devices to Monitor, page 69-34).

    Events from unmonitored devices are always dropped, but they do place a load on the service. The IP address of the last unmonitored device detected is shown; use the IP address to determine the source of the messages. You can then determine if the device should be added to the monitored devices list, or if you need to alter the device's configuration to remove the Security Manager server from its list of syslog servers.

    If the device that is sending messages is outside of your network, adjust the firewall configuration to prevent this syslog traffic from entering your network.
- **Status Information:**
  - **Total Events Per Second (EPS)**—The rate at which events are currently being processed. This measure does not include dropped events.
  - **Event Buffer Used**—The percentage of the shared event buffer that is currently being used to process events. The bar is color-coded to indicate the throttle level:

    Green—Not in throttle mode.

    Yellow—Low severity events are being dropped.

    Orange—Low and medium severity events are being dropped.

    Red—High severity events are being dropped.
- **Event Server Alerts**—These messages indicate specific status problems that you might need to address. Table 69-8 explains the messages that you might see with possible solutions.
- **Copy button**—Click the Copy button to copy the information to the clipboard. The copied information includes HTML markup. You can paste the information into an HTML file.

Table 69-8 Event Manager Status Messages

Alert Message	Alert Level	Possible Action
UDP port <514> could not be acquired, therefore syslog events cannot be collected.	High	Some external application might already be using the indicated port (the default syslog port is 514). You might need to stop that external application. You can use the netstat command to identify the PID of the process, for example, **netstat -ao | findstr 514**.
The event data store location does not exist, therefore events cannot be stored.	High	The event data store location as configured in the Security Manager Administrative Settings does not exist or the Security Manager server does not have the required read/write permissions to the location. For more information about configuring the location, see Event Management Page, page 11-27.
Low severity events are being dropped.	Low	Either events are being received at a very high rate or the system is under a heavy load. To identify if a device is sending events too frequently, you can open the All Device Events view and switch to real-time mode, as described in Switching Between Real-Time and Historical Views, page 69-42. To identify if the server is under a heavy load, log into Windows on the server and use Task Manager or another tool to see if there is an application other than Security Manager that is taxing the system. If possible, disable or stop the application. If the problem occurs frequently, consider uninstalling the other application from the server.
Low and medium severity events are being dropped.	Medium	
All events are being dropped.	High	

Table 69-8 Event Manager Status Messages (continued)

Alert Message	Alert Level	Possible Action
Events from unknown devices are being received.	Low	Syslog events are being sent to the Security Manager server from devices that are not selected for monitoring as described in Selecting Devices to Monitor, page 69-34. These devices might not be supported device types for monitoring and they might not even be in the Security Manager inventory. The message varies based on the EPS rate for these devices. A low severity message indicates the EPS rate is between 500 and 5,000; a medium indicates an EPS rate between 5,000 and 10,000; a high indicates an EPS rate greater than 10,000. The **Events from Unmonitored Devices** statistic in the Last 5 Minutes Statistics shows the number of these events and the IP address of the last unsupported device. Either select the device for monitoring or change the syslog policy for the device to remove the address of the Security Manager server. You will need to repeat the process if more than one unmonitored device is sending messages.
Events from unknown devices are being received at a high rate.	Medium	
Events from unknown devices are being received at a very high rate.	High	

Selecting Devices to Monitor

All ASA and FWSM devices and security contexts, and IPS devices and virtual sensors, that are added to the Security Manager database are automatically selected for monitoring in Event Viewer.

Note To reliably report events from contexts in multiple-context mode, Cisco Event Viewer requires an IP address for the management interface of each context.

Note Beginning with version 4.17, Cisco Security Manager receives events from non-management interfaces also, with following limitations:

- Only Static IP address Interface is supported.
- The events from cluster devices will not be displayed because cluster pool IP ranges for syslog would be used.
- Only after deployment of syslog server configuration, the device event manager is notified to get the non-managed interface IP. Hence, there can be initial event drop.
- This enhancement is not supported in syslog relay services.

If you do not want to use Event Viewer with a device, you can deselect the device for monitoring. Note that if an ASA or FWSM device or security context is not configured to use the Security Manager server as a syslog server, you will not get events from the device or security context anyway, so you might not need to deselect an ASA or FWSM that you do not want to monitor.

Tip You cannot monitor Cisco IOS IPS devices in Event Viewer.

Related Topics

- Adding Devices to the Device Inventory, page 3-7
- Configuring ASA and FWSM Devices for Event Management, page 69-28
- Configuring IPS Devices for Event Management, page 69-29

Step 1 In Event Viewer, select **View > Manage Monitored Device** to open the Manage Monitored Devices dialog box.

The device list shows all devices in the Security Manager inventory to which you have view permissions. You cannot see any devices for which you have no permissions. Any selections you make are limited to the displayed devices. If you do not have permission to select or deselect any devices, the list is read-only and you cannot select devices for monitoring. For more information on access permissions, see Understanding Event Viewer Access Control, page 69-4.

Step 2 Ensure that only those devices whose events you want to monitor in Event Viewer are selected. Deselect any undesired devices.

You can change the selection status for all devices in a device group by selecting or deselecting the group.

Step 3 Click **OK**.

You might need to wait for the changes take effect in Event Viewer.

Monitoring Event Data Store Disk Space Usage

The Event Manager service uses a specified amount of disk space for the primary and extended event data stores. This ensures that the service does not overwhelm the server computer or the extended storage location. You configure the size of the primary and extended event data stores on the **Tools > Security Manager Administration > Event Management** page as described in Event Management Page, page 11-27.

For both the primary and extended locations, when the allocated space is 90% full, the oldest event data is deleted from storage to make room for new data. Data is copied from the primary store to the extended store, if you configure one, so in most cases events deleted from the primary storage continue to be available for querying from the extended storage location, until they are rotated out of the extended storage. (The timing of the copy from the primary to extended data store depends on a number of factors, including the events per second (EPS) rate, the relative size of the primary store to the extended store, and the percentage of the primary data that has already been copied to the extended store.)

You can monitor how much of the allocated space is currently being used, and the age of the oldest event, by selecting **View > Show Event Store Disk Usage** in Event Viewer. The information is displayed as a pie chart that shows the used and unused space in gigabytes (GB) for each location. There is also an indication of the oldest event currently stored in each location.

You can use this information to help you decide whether to increase or decrease the space allocated to each location.

Tip If you decrease the size of either location, and your new size is less than the amount of space currently being used, the oldest events are immediately deleted until your new target size is reached.

Archiving or Backing Up and Restoring the Event Data Store

The event data store is not included with the regular Security Manager database backup. If you want to archive or back up the event data store, whether the primary or extended location, you must do so separately. You can restore the backups if necessary.

This procedure explains the steps required for backup and restore for the event data store.

Tip When you disable the Event Manager service, events are not written to the data store, so you will miss any events generated during the backup or restore process.

Step 1 To back up the event data store:

a. Using the Security Manager client, select **Tools > Security Manager Administration**, and select **Event Management** from the table of contents.

b. Determine the name of the event data store folder. The folder is shown in the Event Data Store Location field; the default is *NMSROOT***MDC\eventing\database**, where NMSROOT is the installation directory (usually C:\Program Files\CSCOpx).

   If you are backing up the extended data store, the location is identified in the Extended Data Store Location field.

c. Deselect the Enable Event Management check box to stop the Event Manager service. Click **Save** to save your changes. You are prompted to verify that you want to stop the service; click **Yes** and wait until you are notified that the service has stopped.

d. Outside of Security Manager, make a copy of the *NMSROOT***MDC\eventing\config\collector.properties** file and the event data store folder. Place the copy on a separate server so that the backup is available in case of hardware failure.

   If you are also backing up the extended data store, make a copy of that folder as well.

e. In the Security Manager client's **Tools > Security Manager Administration > Event Management** page, select the **Enable Event Management** check box and click **Save**. You are prompted to verify that you want to start the service; click **Yes** and wait until you are notified that the service has started.

Step 2 To restore the event data store, use the same process you used to back up the data with the following exceptions:

- Instead of making a copy of the existing event data store, copy the backup into the event data store location. You can optionally delete the existing data before copying in the backup data. However, as long as you do not exceed the data store size limit, you can mix the backup and existing data. (The data store limit is configured in the **Tools > Security Manager Administration > Event Management** page.)

Note Mixing old and new data works only if you are preserving the existing copy of collector.properties (that is, you are not restoring the file), and the new and old data are from the same server. You cannot merge the data store from two or more separate servers.

- Do not restore collector.properties unless you are recovering from a hardware failure or some other event that required you to reinstall Security Manager.

Using Event Viewer

Use Event Viewer to help troubleshoot network problems involving monitored devices. Using views and filtering, you can analyze problems to help identify the cause and possible remedies.

This section contains the following topics:

- Using Event Views, page 69-37
- Filtering and Querying Events, page 69-43
- Performing Operations on Specific Events, page 69-49
- Looking Up a Security Manager Policy from Event Viewer, page 69-54
- Looking Up Events for a Security Manager Policy, page 69-55

Using Event Views

When you view events in Event Viewer, you open a view. A *view* is a set of filters and other properties, including color rules, selected columns and their positions and widths, and the default time window, that let you define a subset of events. Views help to limit the scope of the events list so that you can more easily find what you are looking for.

This section contains the following topics:

- Opening Views, page 69-38
- Floating and Arranging Views, page 69-38
- Customizing the Event Table Appearance, page 69-39
- Switching Between Source/Destination IP Addresses and Host Object Names, page 69-39
- Configuring Color Rules for a View, page 69-40
- Creating Custom Views, page 69-41
- Editing a Custom View Name or Description, page 69-41
- Switching Between Real-Time and Historical Views, page 69-42
- Saving Views, page 69-42
- Deleting Custom Views, page 69-43

Opening Views

You can open up to four historical views and one real-time view in Event Viewer. When you open a view, Event Viewer uses the view settings and time range to retrieve events from the event data store and display them in the event table.

- To open a view so that it replaces the currently active open view, do one of the following in Event Viewer:
  - Double-click the view in the views list.
  - Right-click the view in the view list and select **Open**.
- To open a view in a new tab, do one of the following:
  - Select **File > Open View** from the menu bar. The Open a View dialog box opens, which is essentially the same as the view list. Select the view and click **OK**.
  - Right-click the view in the views list and select **Open In New Tab**.

Floating and Arranging Views

You can open up to four historical views and one real-time view at one time. When you have multiple views open, they are opened as tabbed windows in the right pane of the main Event Viewer window, in the most recently used area ("tabbed group") if there is more than one area. The commands to arrange the windows appear if you right-click the tab for the view window as shown in the following illustration.

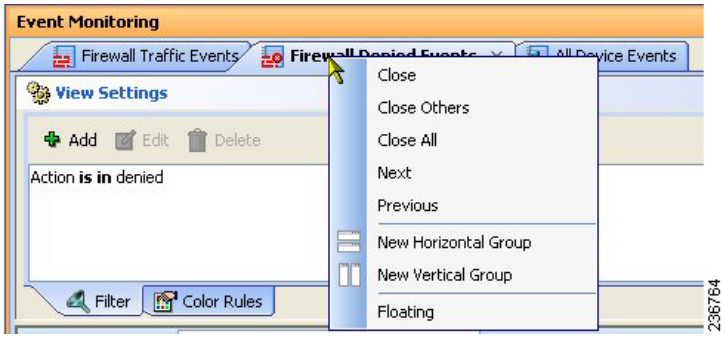

You have many options for arranging view windows based on your requirements. For example, you might want to compare two views side-by-side, or remove a view from the main window without closing it.

You can use the following techniques to arrange the view windows to get the display that you desire:

- Floating a view—To remove a view from the main Event Viewer window without closing it, right-click the view tab and select **Floating**. The view is moved to its own window.

  If you have already floated a view, you can select **Floating to** and choose one of the already-floated windows. The view becomes a new tab in that window.
- Docking a view—To move a floating view back to the main Event Viewer window, right-click the view tab and select **Docking**.
- Arranging views horizontally or vertically for side-to-side comparison—To create a vertical or horizontal arrangement of views to allow for easy comparison, without floating the views, right-click the view tab and select **New Horizontal Group** or **New Vertical Group**. These

commands split the current tabbed group into the selected layout. You must have at least two views open to use these commands. If you have more than two open views, and you want all of them in separate windows, you need to use the command multiple times.

- Move views to different tabbed groups—If you have several open views, and you have arranged them into horizontal or vertical groups, you can move views among the groups by right-clicking the view tab and selecting **Move to Next Tab Group** or **Move to Previous Tab Group**. The commands appear only if views are arranged in a manner where such movement is possible:
- Change the orientation of groups—You can switch between horizontal and vertical layouts by right-clicking the view tab and selecting **Change Tab Groups Orientation**.

Customizing the Event Table Appearance

You can customize the appearance of predefined or custom views in the event table to meet your requirements. You can save these changes even in the predefined views.

You can do the following to customize the event table:

- Create column filters to limit the type of events listed. Use the down arrows in the column heading to define the filter as described in Creating Column-Based Filters, page 69-45.
- Create color rules to highlight events based on severity, as described in Configuring Color Rules for a View, page 69-40.
- Change which columns appear in the table by clicking the Column Selector icon to the right of the table heading row, as described in Columns in Event Table, page 69-18.
- Change the width of a column by clicking the right edge of the column heading and dragging it to the desired size.
- Change the order of the columns by clicking the column heading and dragging the column to the position you want.
- Sort the events list by a column by clicking the column heading. The column sorts based on a three-click cycle: ascending, descending, and default order (which is by event reception time).
- Reset the width of the View Selector and Event Monitoring window to their default values by selecting **View > Reset Layout**.
- Change whether the source and destination columns show IP addresses or host object names, as described in Switching Between Source/Destination IP Addresses and Host Object Names, page 69-39.

Related Topics:

- Creating Custom Views, page 69-41
- Saving Views, page 69-42

Switching Between Source/Destination IP Addresses and Host Object Names

You can view source and destination IP addresses or you can view the host object name of objects that match a source or destination IP address. By default, the Event Viewer shows host object names when available.

IP address to host name mapping is supported only for the source and destination of events. Also, the mapping applies to Host objects only; Event Viewer will not show an object name when the source or destination of an event matches a Network object, Group object, or Address Range object.

Please clarify the type and content of the object. E.g. is this feature for host type network/host objects only, that is, the single-value host version of the object? does it work for single-value group objects, or for network or range objects?

To switch between source/destination IP addresses and host object names, do the following:

- To see host objects names for any objects that match a source or destination IP address, select **View > Show Network Host Objects**. This option is selected by default.

Tip Hover over a host object name to view the IP address associated with that object.

Note An IP address to host object name cache is created when Event Viewer is launched. If you define new host objects, you must submit those changes to the database and then close and relaunch Event Viewer for those mappings to be used.

- To see IP addresses in the source and destination columns, deselect **View > Show Network Host Objects**.

Configuring Color Rules for a View

You can use color rules to color-code events shown in the event table based on the severity of the event. Color-coding can help you quickly identify the events that you most want to know about.

You can selectively enable and disable color rules by editing them. This allows you to turn them on and off without deleting them.

Tip You can configure color rules for both predefined and custom views. However, you cannot share color rules between views: all color rules are unique to a view. If you want to apply the same rules to multiple views, you must recreate the rules in each view.

To define and enable a color rule, follow these steps:

Step 1 Open the view in which you want to define the color rule (see Opening Views, page 69-38).

Step 2 Click the **Color Rules** tab in the View Settings pane (see Event Monitoring Window, page 69-14).

Step 3 Do any of the following:

- To add a new rule, click the **Add** button. In the Add Color Rule dialog box, configure the rule as follows:
  - Select **Enable** to make the rule active.
  - Select the severity level for which the rule applies from the **Severity** list.
  - Use the **Foreground** (which is the text color), **Background**, and **Font Type** (either Bold or Italics) controls to define how the severity should be presented in the table. The Preview Text area shows how your rule will look.
- To edit a rule, select it and click the **Edit** button.
- To delete a rule, select it and click the **Delete** button.

Creating Custom Views

A custom view is one in which you define the filters in the view settings. Using custom views, you can configure filter rules to pin-point specific areas for monitoring and analysis. Custom views are private and cannot be shared between users.

You basically have two options for creating custom views, creating a view from scratch or from an existing view:

- To create a custom view that has no predefined column filters, do one of the following:
  - Select **File > New View** from the menu bar.
  - Click the **New** button above the view list.

  Then, enter a name for the view and optionally a description of the view and click **OK**. The view is added to the My Views folder in the views list.
- To create a custom view based on an existing view, do one of the following:
  - With the desired base view open, click the down arrow on the Save button in the event table toolbar and select **Save As**, or select **File > Save As** from the menu bar.
  - Right-click the desired base view in the views list and select **Save As**.

  Then, enter a name for the view and optionally a description of the view and click **OK**. The view is added to the My Views folder in the views list. The new view has the same filters as the base view.

Note View names can be up to 128 characters and contain alphanumeric characters, spaces, hyphens (-), underscore characters (_), plus signs (+), periods, and ampersands (&). The description can be up to 1024 characters.

After you create the new view, you can customize it the same way that you can an existing view:

- Define filters in the view settings. See Creating Column-Based Filters, page 69-45.
- Define color rules in the view settings. See Configuring Color Rules for a View, page 69-40.
- Select the columns to display in the event table. See Columns in Event Table, page 69-18.
- Customize the event table appearance. See Customizing the Event Table Appearance, page 69-39.

Related Topics

- Views and Filters, page 69-3
- Event Table Toolbar, page 69-16
- Editing a Custom View Name or Description, page 69-41
- Deleting Custom Views, page 69-43

Editing a Custom View Name or Description

To change the name of a custom view, or the custom view's description, do one of the following:

- Select the custom view in the view list and click the **Edit** button above the list.
- Right-click the custom view in the view list and select **Edit**.

Then, make the desired changes to the custom view name or description and click **OK**.

Note View names can be up to 128 characters and contain alphanumeric characters, spaces, hyphens (-), underscore characters (_), plus signs (+), periods, and ampersands (&). The description can be up to 1024 characters.

You cannot change the name or description of a predefined view.

Switching Between Real-Time and Historical Views

You can update the events table for any view using either real-time or historical time periods. A real-time view shows events as they are received, whereas an historical view shows a static list of events that is not updated until you click the **Start** button in the event table toolbar.

To switch between real-time and historical time periods in an open view, do the following:

- To see events in real-time, select **View > Mode > Real Time**, or click the Time Selector control in the event table toolbar and select **Real Time**. For help in locating the control on the toolbar, see Event Table Toolbar, page 69-16.
- To see events in an historical period, select the desired time frame from the **View > Mode** menu or from the Time Selector control on the event table toolbar. All options other than Real Time are historical views. For more information, see Selecting the Time Range for Events, page 69-43.

Saving Views

If you edit the settings for a view, you must save it to make those changes permanent. Saving a view saves changes to filters (for custom views only), table preferences such as selected columns, column width, and sort order, the time range, and color rules. If you make filter changes to a predefined view, you must use Save As to create a new custom view.

- To save changes to a view, do one of the following in Event Viewer:
  - Select **File > Save** from the menu bar.
  - Click the **Save** button in the event table toolbar.

  You are asked to confirm that you want to save your changes.
- To save your changes as a new custom view, do one of the following to open the Save View As dialog box:
  - Select **File > Save As** from the menu bar.
  - Click the down arrow on the Save button in the event table toolbar and select **Save As**.
  - Right-click the view in the views list and select **Save As**.

  Then, enter a name for the view and optionally a description of the view and click **OK**. The view is added to the My Views folder in the views list.

  Note View names can be up to 128 characters and contain alphanumeric characters, spaces, hyphens (-), underscore characters (_), plus signs (+), periods, and ampersands (&). The description can be up to 1024 characters.

Deleting Custom Views

You can delete custom views, but you cannot delete predefined views. To delete a custom view, do one of the following:

- Select it in the view list and click the **Delete** (trash can) button above the list.
- Right-click it in the view list and select **Delete**.

You are asked to confirm your deletion.

Filtering and Querying Events

There are many options for filtering the events that appear in the event table. You can reduce the list of events by selecting the appropriate time range, by filtering on elements in specific columns, or even by searching on a text string.

This section contains the following topics:

- Selecting the Time Range for Events, page 69-43
- Using the Time Slider with Filtering, page 69-44
- Refreshing the Event Table, page 69-44
- Creating Column-Based Filters, page 69-45
- Filtering Based on a Specific Event's Values, page 69-47
- Filtering on a Text String, page 69-47
- Clearing Filters, page 69-48

Selecting the Time Range for Events

Use the Time Selector control in the event table toolbar, or the equivalent **View > Mode** command, to select the time range for displaying events. The event table lists only those events that occur within the selected time range. For help in locating the Time Selector control on the toolbar, see Event Table Toolbar, page 69-16.

Tip For historical views, the time is based on server time, not the time configured on your workstation.

When you change the time range, the table reloads to show events in the selected range. For historical views, you can refresh the events list by clicking **Start** or by doing the other actions described in Refreshing the Event Table, page 69-44.

The following are your options for the time range:

- To view events from the present time into the past, select one of the following time periods: **last 10 minutes**, **last 1 hour**, **last 12 hours**, **last 1 day**, or **last 1 week**.
- To view events from today or yesterday, select **today** or **yesterday**, as desired.
- To view events from a specific day, select **is on** and then select the date from the displayed calendar.
- To view events from a specific date and time range, select **is between** and select the first and last days and times from the displayed calendars.
- To view real-time events, select Real Time.

Using the Time Slider with Filtering

You can use the vertical slider control in the time slider to change the start time for the events shown in the event table. This is particularly useful when you want to locate events and you know the approximate time they occurred.

For details on the operation of the time slider, see Time Slider, page 69-25.

To use the time slider to aid filtering, follow these steps:

Step 1 Open an historical view, or use the Time Selector in the toolbar, or the **View > Mode** command, to select an appropriate time range, such as Last 10 Minutes. For more information, see Selecting the Time Range for Events, page 69-43.

Step 2 Move the vertical slider to the approximate time of the events you want to examine.

The event table is reloaded to display events on or before the time you specified with the vertical slider. This time range is shaded in the time slider, and the length of time selected is based on the length of time selected in the Time Selector (for example, 10 minutes for a Last 10 Minute view, or the same length of time selected for an "is between" view). The time range is noted in the Time Selector control, for example:

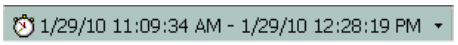

Step 3 To locate the events, you can now do any of the following:

- Apply custom column filters. See Creating Column-Based Filters, page 69-45.
- Use the quick filter to search on a text string. See Filtering on a Text String, page 69-47.
- Scroll or page through the event table.
- Use the time slider paging controls to reset the time range forward or back. For more information, see Time Slider, page 69-25.

Note The distance moved forward or back when paging in the time slider depends either on the mode (time range) that is set or the number of events the event table can hold. The position of the vertical slider denotes the most recent event loaded in the event table.

Refreshing the Event Table

When you are using an historical mode, such as "last 10 minutes," the latest events displayed correspond to the time you selected the time range or opened the view. Similarly, if you are in real-time mode and have clicked Stop, the event table does not include events that arrived after you stopped the event stream.

To refresh the events listed in the event table, so that they are current with your selected time range, do any of the following:

- Click **Start** in the toolbar, or select **View > Start**. The table is refreshed based on your currently selected time range. For real-time views, the event stream restarts.
- Select a different time range using the Time Selector in the toolbar or the **View > Mode** command.
- Select a different time slice using the vertical slider or the pagination controls in the time slider below the event table. For more information on using these controls, see Time Slider, page 69-25.

Creating Column-Based Filters

You can filter the event table in Event Viewer based on the contents of specific columns. Column filters are the type of filter contained in the view settings; they define the basic content of the view. Whenever you apply a column filter, the view settings for the view are updated to include the newly selected filter: you must save the view before closing it if you want the new filter to become a permanent part of the view's definition.

There are many ways in which to define a column filter:

- In the View Settings pane, click the **Add** button. You are first prompted to select the column on which to base the filter. When you click **OK**, you are prompted to create the filter.
- In the View Settings pane, select a filter and click the **Edit** button to change it.
- In the event table, click the down arrow button in the heading of a column and select any of the following from the drop-down list:
  - A specific entry. The drop-down list contains all values currently displayed in the events listed in the table.
  - (All). Select (All) to remove a filter from this column. The event table is updated to show the events that meet your other filter criteria.
  - (Custom). Select (Custom) to create a filter that might have multiple values, negative values, or be based on data not currently contained in the column in the current event table. Selecting (Custom) is essentially the same as creating a filter directly in the View Settings pane.
- In the event table, you can right-click a value and select **Filter This Value**. This action has the same effect as selecting the value from the drop-down list for the column.

  You can alternatively select **Filter Not This Value** to create a filter that excludes a value,
- In the event table, you can right-click a value and select **Create Filter from Event**. You are prompted to select the specific columns to include; the column on which you right-clicked is initially selected, but you can deselect it.

The following procedure explains how to build a custom column-based filter, one in which you are not simply selecting a value from the column's drop-down list.

Tips

- Column filters are cumulative: for an event to appear in the event table for a view, the event must meet all column filter criteria. You cannot create a set of OR'ed column filters.
- Some columns allow you to select network/host or service policy objects to define the filter criteria. Selecting policy objects can simplify your filters. However, for a policy object to be selectable in a filter, the object must be committed to the database. If you create a new object for filtering purposes, ensure that you submit your changes in Configuration Manager (and if using Workflow mode with an approver, get the changes approved) before attempting to create the filter in Event Viewer.

  When using policy objects, the filtering recognizes whether a device-level override is defined for the object. For example, if you use a network/host object that contains 10.10.10.10, and Device A has an override to change the address to 10.10.10.12, events from Device A appear in the list only if the event matches 10.10.10.12. For devices that do not have overrides, the events must match 10.10.10.10. Furthermore, if Device A has an event that matches 10.10.10.10, that event is not listed because it does not match the device-level override. Thus, using policy objects can provide results that vary by device and therefore match more closely to your policy definitions.

  If your organization is using ACS to control user access, you must have the appropriate View Object privileges for network/host, network/host-IPv6, and service objects to use them in filters.

- You can filter on the contents of most but not all columns. If a column does not have a down arrow, you cannot filter on it. For example, you cannot filter on Description, Event Name, Generation Time, or Receive Time.
- The filter icon (a funnel) appears in the heading of a filtered column.
- For a description of the available columns, see Columns in Event Table, page 69-18.

Step 1 Do one of the following:

- In the View Settings pane, click the **Add** button. The Add Custom Filter to a Column dialog box opens. Select the column on which to base the filter and click **OK**.
- In the View Settings pane, select a filter you want to change and click the **Edit** button.
- From the drop-down list for a column, select **(Custom)**.
- Right-click any cell in the desired column and select **Custom Filter**.

The Custom Filter dialog box opens for the selected column.

Step 2 In the Custom Filter dialog box, select the desired values. The following illustration shows a typical example of this dialog box, for the Source column.

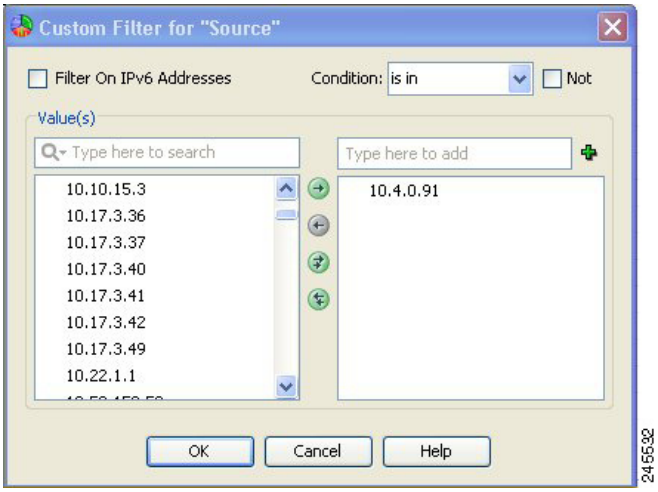

Following is a description of the controls you might find in the Custom Filter dialog box. Not all controls appear for every column:

- **Available and Selected Items lists**—In most cases, to select an item you highlight it in the left list, which contains the available values, and click the right arrow to move it to the list of selected values. You can select multiple values. The right list defines the filtered values.

  The items listed in the available values column are determined by the values currently present in the events listed in the events table. For address and service fields, the list also includes policy objects. If there are a lot of available values, you can search for the desired value by typing into the edit box above the list; the list is filtered as you type. Click the down arrow next to the Q to change how your search string is evaluated for matches.

  You can also select, or deselect, values using the following techniques:

  – Type the item into the edit box above the selected values list and click the + button. This technique is useful if there is a large number of available values, or if you want to filter on a value that is not present in the current events list.

- Double-click an item in either list to move it to the other list.
- Click the double-arrow buttons to move all items, regardless of your selection.

Note In a limited number of cases, the Custom Filter dialog box contains a single list. For example, the dialog boxes for the Event Type ID and Device columns contain single selectors. In these cases, make your selection using the check boxes next to the items; selecting a folder selects all items in the folder.

- **Filter on IPv6 Addresses**—For columns that contain addresses, use this option to toggle between listing IPv4 and IPv6 addresses and network/host objects in the available values column. You can filter on either IPv4 addresses or IPv6 addresses, but not both, in a single view.
- **Condition, Not**—Defines the condition applied to the selected items, typically "is in."

  To create a negative condition, so that selected values define the events to not include in the events table, select the Not option.

Step 3 Click **OK**.

The view settings are updated to include the new filter, and the events table is updated to show only those events that satisfy all filters.

Filtering Based on a Specific Event's Values

You can base a new filter on information contained within an event, or a single cell within an event, by right-clicking and choosing a filter command. When you filter using these commands, a column filter is added to the view settings. You can do the following:

- To create a filter based on multiple values in the selected event, select **Create Filter from Event**, then select from the dialog box the values on which to filter. The dialog box lists only those columns that are displayed in the table; the current values are shown in parentheses. For an explanation of the columns, see Columns in Event Table, page 69-18.
- To filter on only the value in the cell on which you right-click, select **Filter This Value**.
- To filter to exclude the value in the cell on which you right-click, select **Filter Not this Value**. All events that do not contain the selected value in this column, including all empty cells, are shown in the table.
- To filter on the flow of the selected event, based on source, source service, destination, and destination service, select **Filter This Flow**.

Filtering on a Text String

Use the quick filter to search for text strings in events. As you type a search keyword, the events table automatically excludes non-matching events as you type. You can search on all columns (the default), or you can select a specific column in which to search.

The following illustration shows the quick filter, which is on the right of the event table toolbar (see Event Table Toolbar, page 69-16).

To perform a search, simply type in the search string. To change how the string is evaluated, click the down arrow next to the Q (magnifying glass) in the left of the edit box. You can limit the search scope using these controls:

- Column name—Select a specific column to search only within that column. The list includes all columns currently displayed in the table. The default is to search all columns.
- Case sensitivity—Select **Case sensitive** or **Case insensitive** to control whether capitalization is considered when selecting matches. The default is case insensitive.
- Wild card usage—Select **Use Wild Cards** to have the following characters evaluated as wild cards:
  - * (asterisk)—matches 0 or more characters.
  - ? (question mark)—matches one character.
- Match method—Select one of the following to determine the location within a cell that the search string should reside:
  - Match from start—The string must be at the beginning of the cell.
  - Match exactly—The cell must contain all and only the search string.
  - Match anywhere—The string can appear anywhere within the cell.

To remove the search string, simply delete it from the quick filter edit box.

For example, if you want to find events that relate to ports that start with tcp/48, type **tcp/48** into the quick filter. In the following illustration, note that all but six events are filtered out of the table. In this example, the search string is found in the Source Service column for the first five events, but in the Destination Service column for the sixth event. If you know beforehand that you are interested in destination services only, you could select **Destination Service** from the quick filter drop-down list and the table would show the last event only.

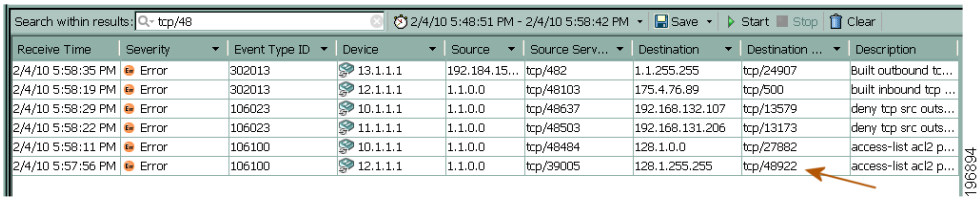

Clearing Filters

When you apply filters to the event table, non-matching events are not displayed. You might find that you need to see the non-matching events. You can either open a different view that applies different (or no) filters, or you can clear filters from the current view.

When clearing filters, the filter definition is removed from the view settings, but the change is not permanent until you click **Save**. Thus, you can remove filters temporarily without redefining the view settings.

You can clear filters one at a time or clear all filters:

- To clear a single filter, so any of the following:
  - Select the filter in the View Settings pane and click **Delete**.
  - Select **(All)** from the drop-down list of a filtered column.
  - Right-click in the filtered column and select **Clear This Filter**.

- To clear all filters, right click in the events table and select **Clear all filters**.

Performing Operations on Specific Events

You can operate upon a single event in the event table in a variety of ways, which include the following:

- **Right-click**—Right-clicking a single event in the event table opens a context menu with commands that you can use on the event. For more information about what you can do from the right-click menu, see the following topics:
  - Event Context (Right-Click) Menu, page 69-49
  - Examining Details of a Single Event, page 69-53
  - Copying Event Records, page 69-53
  - Saving Events to a File, page 69-53
  - Filtering Based on a Specific Event's Values, page 69-47

  Note You can hover your mouse over a valid IPv4 address in Event Viewer to launch the IP Intelligence tool for that IP address. The IP Intelligence tool provides various pieces of information about an IPv4 address, such as the fully qualified domain name (FQDN), geographic location information, and WHOIS information. For more information on the IP Intelligence tool, see IP Intelligence, page 72-34.

- **Select an event**—When you click a single event in the event table it is highlighted and the Event Details pane displays details for that particular event. Hold the **Ctrl** key to select additional events, or hold the **Shift** key to select a range of events.
- **Double-click an event**—Double-clicking a single event in the event table opens the Event Details dialog box, which shows the event information in an easier-to-read format. From the Event Details dialog box, you can print the displayed details or copy some, or all, of the details to the clipboard for pasting into another program. You can use the Next and Previous buttons to scroll through the events listed in the event table. For information on the meaning of the attributes, see Columns in Event Table, page 69-18.

  Alternatively, you can right-click on an event and select **Show All Details** to open the Event Details dialog box.

Event Context (Right-Click) Menu

When you right-click an event in the Event Table, a context menu appears that provides commands that you can use with the selected event. The specific list of available commands depends on the type of event and also the specific cell on which you right-click. The following table explains all of the available commands.

Note In addition to the right-click options listed below, you can also hover your mouse over a valid IPv4 address in Event Viewer to launch the IP Intelligence tool for that IP address. The IP Intelligence tool provides various pieces of information about an IPv4 address, such as the fully qualified domain name (FQDN), geographic location information, and WHOIS information. For more information on the IP Intelligence tool, see IP Intelligence, page 72-34.

Table 69-9 Event Context Menu

Command	Description
Clear This Filter	Removes the filter defined for this column. The command is available only if you right-click on a cell that is in a filtered column. The filter is removed from the view settings. You must save the view to make your change permanent.
Clear All Filters	Removes all filters from the view settings. This command is available only if there is at least one column filter. You must save the view to make your change permanent.
Filter This Value Filter Not This Value	Creates a column filter based on the value in the cell that you right click. You can create either a positive or negative filter based on the value. The view settings are updated with the new filter and replace any existing filter for this column. You must save the view to make your change permanent.
Create Filter from Event	Creates a set of column filters based on the values in the selected event. You are prompted to select the specific columns to include; the column on which you right-clicked is initially selected, but you can deselect it. The view settings are updated with the new filters and any existing column filters for the selected columns are replaced. You must save the view to make your changes permanent.
Custom Filter	Creates a custom column filter, as described in Creating Column-Based Filters, page 69-45. The view settings are updated with the new filter and any existing filters for the selected columns are replaced. You must save the view to make your changes permanent.
Filter This Flow	Creates a set of column-based filters that present the events related to a specific traffic flow. Filtered columns are source and source service and destination and destination service. The view settings are updated with the new filter and any existing filters for the selected columns are replaced. You must save the view to make your changes permanent.
Show IPLogs	Opens the IP log for an IPS Alert event using an external packet analyzer tool. You must have a packet analyzer installed and associated with *.pcap file extension.
Show All Details	Opens the Event Details dialog box for the event, which shows all event information in an easier-to-read format. You can also print the details or copy them to the clipboard. The details are the same as those shown in the Event Details pane below the event table.

Table 69-9 Event Context Menu (continued)

Command	Description
Copy commands	You can use the following commands to copy event data to the clipboard. You can then paste the data into a spreadsheet or other program for your use. For more information, see Copying Event Records, page 69-53. • **Copy Cell**—Copies the contents of the cell you right-click to the clipboard. • **Copy Selected Events**—Copies the contents of all selected (high-lighted) events to the clipboard. • **Copy All Events**—Copies the contents of all listed events to the clipboard. This command is useful only if you have filtered the event table to a manageable number of events.
Save Selected Events as HTML Save All Events as HTML Save Selected Events as CSV Save All Events as CSV	Saves either all events listed in the event table, or all selected (high-lighted) events, to an HTML or comma-separate values (CSV) file on your workstation. You are prompted to select the folder and enter the file name for the export file. For more information, see Saving Events to a File, page 69-53.
Go To Policy	Finds the policy that generated this event in the device's policy configuration in Configuration Manager. This command is available only for events where a binoculars icon appears in the Event Name cell. For detailed information, see Looking Up a Security Manager Policy from Event Viewer, page 69-54.
Packet Capture	Opens the packet capture tool, where you can define criteria for capturing packets on the device.
Ping and TraceRoute	Opens the Ping, TraceRoute, and NS Lookup tool, where you can use these applications with the device from which the event was sent. For detailed information, see Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools, page 72-26

Table 69-9 Event Context Menu (continued)

Command	Description
Tune Signature	Opens the IPS Signature Quick Tune dialog box where you can enable or disable the signature associated with the selected event, and modify the Base Risk Rating of the signature that is assigned to the device or shared policy. To tune a signature you must create or open a ticket. For more information see Working with Activities/Tickets, page 4-7. The Base Risk Rating value of the signature is calculated by multiplying the fidelity rating and the severity factor and dividing them by 100 (Fidelity Rating x Severity Factor /100). This value is read only; you cannot directly change it. To change the Base Risk Rating, you must alter the Severity and Fidelity values. • Severity: The severity level that the signature will report: High, Medium, Low, or Informational, where, – High = 100 – Medium = 75 – Low = 50 – Informational = 25 • Fidelity: The Fidelity Rating, or Signature Fidelity Rating (SFR), identifies the weight associated with how well this signature might perform in the absence of specific knowledge of the target. This rating can be any number from 0 to 100, with 100 indicating the most confidence in the signature. After you enable or disable the signature or modify the Base Risk Rating you must redeploy the configuration to the device, using Configuration Manager, for the change to take effect on the device. Note that such changes will affect only real time events and not the historical events. For information about deploying configuration, see Chapter 8, "Managing Deployment".

IPS Signature Quick Tune Dialog Box

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Use the IPS Signature Quick Tune dialog box to enable or disable the signature associated with the selected event, and modify the Base Risk Rating of the signature that is assigned to the device or shared policy.

Navigation Path

In Event Viewer, right-click a row (an event) and click **Tune Signature**. For more information, see Event Context (Right-Click) Menu, page 69-49.

Examining Details of a Single Event

Each event contains a lot of specific information in many separate fields. Typically, you display a subset of these fields in the event table. When you want to see the complete details of an event, you can use either of the following:

- **Event Details pane**—Select the event and open the Event Details pane below the event table. You can open the pane by clicking anywhere in the "Event Details" title row, or you can select **View > Show Event Details** from the menu. The Event Details pane organizes the information in tabs. For more information about this pane, see Event Details Pane, page 69-26.
- **Event Details dialog box**—You can open this dialog box by double-clicking the event, or by right-clicking the event and selecting **Show All Details**. The information is presented as a flat list and shows the information that would be shown on the Details tab in the Event Details pane. For information on the meaning of the attributes, see Columns in Event Table, page 69-18.

  The Event Details dialog box includes the following controls:

  - Print button—Click this button to print the information. You are prompted to select a printer.
  - Copy button—Click the down arrow on this button and select **All Rows** or **Selected Rows**. The information is copied to the clipboard, and you can paste it into another application. Note that the Selected Rows command works only if you select at least one row in the table.
- Next, Previous buttons—Click these buttons to scroll through the events currently displayed in the event table. Next moves up and Previous moves down in the table.

Copying Event Records

You can copy single events, multiple events, all events, or even the contents of a single cell to the clipboard. You can then paste the information into another application, such as a spreadsheet or an e-mail message.

You can do the following from the event table:

- **Copy selected events**—To copy one or more selected events, right-click in the event table and select **Copy Selected Events**. The event you right-click does not matter, the copied events are those that are selected (highlighted) in the table.

  Click an event to select it. Hold **Ctrl** key to select additional events, or hold the **Shift** key to select a range of events.
- **Copy the contents of a single cell**—To copy the contents of a single cell in one event, right-click the cell and select **Copy Cell**. You cannot copy cell contents if there is more than one event selected in the table.
- **Copy all events**—To copy all the events shown in the event table, right-click anywhere in the table and select **Copy All Events**.

Saving Events to a File

Rather than copying events to the clipboard and pasting them into another application, you can directly save events to an HTML or comma-separated values (CSV) file. HTML files are useful for viewing information, whereas you can open a CSV file in a spreadsheet application for further analysis and report generation.

When saving event data, you are prompted to select a folder and enter a file name.

You can do the following from the event table:

- **Save selected events**—To save one or more selected events, right-click in the event table and select either **Save Selected Events as HTML** or **Save Selected Events as CSV**. The event you right-click does not matter, the saved events are those that are selected (highlighted) in the table.

  Click an event to select it. Hold **Ctrl** key to select additional events, or hold the **Shift** key to select a range of events.
- **Save all events**—To save all the events shown in the event table, right-click anywhere in the table and select either **Save All Events as HTML** or **Save All Events as CSV**.

Looking Up a Security Manager Policy from Event Viewer

In Event Viewer, if an event was generated from an IPS signature policy, or from certain actions related to explicit access rules (such as denied access), you can quickly locate the related signature or access rule from the event itself.

The main reason you would want to perform policy lookup is to adjust a policy based on the events that it is generating. For example, an access rule might be dropping traffic that you actually want to allow. Because you are looking at the event, you know there is a policy that is causing the event, so with a few clicks, you can get from that event to the policy that you need to reconfigure.

You can look up policies from the following types of events:

- Firewall events—You can look up policies for the following syslog messages:
  - 106023—Denied IP packet.
  - 106100—Permit/Denied by ACL.
  - 302013—Built TCP (started a TCP session).
  - 302015—Built UDP (started a UDP session).
- IPS alert events—All IPS events that have valid signature and sub-signature identifiers.

Tips and Caveats

- You cannot look up firewall policies for events that contain IPv6 addresses. You can look up IPS policies for IPv6 addresses, however.
- When a policy that is based on IP address alone and not on a user name triggers an event, the device looks up the IP address in the Active Directory and if a user name is associated with that IP address, the user name is added to the syslog. Hence, even if a policy does not contain a user name, the resulting syslog might contain it. Policies cannot be created with a destination user and, as a result, this field will not be used during policy lookup.
- If an event is generated for a policy that is configured based on source/destination FQDN, the resulting syslog will not contain the FQDN because of a device defect. In such cases, policy lookup will not work.
- If an event is generated for a policy that is based on user groups, the syslog will contain the specific user name that triggered the event and not the user group. In such cases, policy lookup will not work.
- Hash codes are required for successful policy lookups from syslog 106023 and 106100 events. These hash codes are available only if you deployed the configuration using Security Manager. If policy lookup fails, try deploying the configuration (either to the device or to a file), then try the policy lookup again.
- If you had applied a filter to the device's policy table, and the rule or signature that generated an event is filtered from the current view, Security Manager cannot highlight it. Clear the filter and try again.

- If the event is caused by an implicit rule, such as the implicit **deny any** at the end of access rules, Security Manager cannot highlight the rule. It is considered good practice to create an explicit deny any rule at the end of access lists.
- The target policy is always found in Device view, even if the device uses a shared policy. Device view is opened if necessary to highlight the policy.
- For IPS signatures, you might not be able to edit the signature if it is a default signature.
- For access rules, the selected rule is the best match for the event. It is possible that more than one rule would generate the same event if you have overlapping or redundant rules. In these cases, editing the selected rule might not completely eliminate the event, because a subsequent rule might perform the same action. Use the access rules tools to analyze and combine overlapping rules.
- For access rules, multiple rules might permit a packet during session creation, but the first rule only is highlighted.
- If your organization is using ACS to control access, you must have View Device privileges to the device, and also View privileges to the firewall or IPS policy, to perform policy lookup. If you do not have all permissions, you will get an "Unable to Find Matching Rule" error if you try to look up a matching rule.

Step 1 Right-click the event in Event Viewer and select **Go To Policy**.

Tip You can identify whether you can look up policies from the event by looking at the Event Name cell in the table. If there is a binoculars icon before the event name, policy lookup is available. Also, if the Go To Policy command is greyed out, you cannot look up policies for that type of event.

Step 2 Security Manager finds the related access rule or IPS signature for the device and highlights it in the policy table. From here, you can edit the policy to view or change it; for detailed instructions, see Configuring Access Rules, page 16-7 and Configuring Signatures, page 39-4.

Your changes do not take effect until you submit and deploy the updated configurations.

Looking Up Events for a Security Manager Policy

You can look up events in Event Viewer that relate to specific firewall access rules or IPS signatures. You can also look up events that relate to specific devices or site-to-site tunnels in Health and Performance Monitor.

When Event Viewer receives events, they are parsed, "sessionized," written to an event buffer, and then written to the database. Sessionizing takes two forms: with a session-oriented protocol, such as TCP, the session encompasses the initial handshake to the connection tear-down; with a sessionless protocol, such as UDP, the session start and end times are based more on first and last packets tracked within a restricted time period—packets that fall outside of the time period are considered parts of other sessions.

Because there is a difference between newly-received and fully processed data, you can look up either real-time or historical events:

- **Real-time**—Because sessionization takes time, keeping an event in cache for up to two minutes, you can use the real-time event query to view events right after parsing, providing access to the most current data received.

- **Historical**—Historical event reports help you identify trends over longer periods of time than is possible with real-time monitoring. For historical events, the Result Format is the All Matching Events option, and the Filter By Time value is set to the previous 10 minutes.

The following topics explain event lookup in more detail:

- Viewing Events for an Access Rule, page 69-56
- Viewing Events for an IPS Signature, page 69-57
- Viewing Events for HPM Devices and Site-to-Site VPNs, page 69-58

Viewing Events for an Access Rule

From the **Firewall > Access Rules** policy in Security Manager, you can select an access rule and view related event information in Event Viewer. You can view real-time or historical events matching the rule. You can view events for ASA (including ASA-SM) and FWSM devices.

Firewall access rules are presented in the form of an ordered list or table. When deployed, this policy becomes an access-control list (ACL), with each entry in the list known as an access-control entry (ACE). (For more detailed information, see Understanding Access Rules, page 16-1.)

When deciding whether to forward or drop a packet, a device tests the packet against each access rule in the ordered listed. If you enable logging for an access rule, the results of the test are recording according to your per-rule log settings. Some devices, such as ASA, generate log entries for denied access even if you do not configure logging explicitly. For information on creating access rules, including logging options, see Configuring Access Rules, page 16-7.

If logging is enabled for the rule (in the Advanced and Edit Options Dialog Boxes, page 16-17), the device sends syslog messages to Event Viewer to record the logged events. This query includes the access-rule parameters, including available keyword information. Reported events do not include connection set-up and tear-down.

To view rule-related events, use the following right-click commands:

- **Show Events > Realtime**—To view real-time query results in Event Viewer for events matching this rule. You can change the query criteria in the Event Monitoring window at any time, applying new parameters to alter the real-time results.
- **Show Events > Historical**—To view historical query results in Event Viewer for events matching this rule. You can change the query criteria in the Event Monitoring window at any time, applying new parameters to alter the historical results.

Security Manager provides the following information to Event Viewer as criteria for access-rule event queries:

- Device details—General information about the device, such as host name, domain name, management IP address, and display name.
- Source addresses—Source addresses of hosts and the network/host objects expanded to display the networks or collections of IP addresses.
- Destination addresses—Destination addresses of hosts and the network/host objects expanded to display the networks or collections of IP addresses.
- Service—Protocol and port information.
- Event Type—“Built/teardown/permitted IP connection” for permit rules and “Deny packet due to security policy” for deny rules.

Notes:

- You can query on only one access rule at a time.
- When NAT or PAT is configured on a security device, the source and destination addresses are mapped to pre-translation and post-translation addresses, respectively, and the translated addresses are used when Security Manager sends a query to Event Viewer. For inbound access rules, the destination address is considered the pre-translation address, and for outbound access rules, the source address is considered the post-translation address.
- Filtering with multiple services (like UDP, TCP, and ICMP) might not give accurate results. To work around this problem, you can remove some of the filters after Event Viewer is launched.
- Filtering based on ICMP sub types is not supported. For example, if an ACE has 'ICMP Echo' in service, the filter is applied only for the protocol (ICMP), but not for type column (Echo) in Event Viewer.
- Service ports with 'eq', 'neq', 'gt', and 'lt' are not supported in cross launch to Event Viewer.

Related Topics

- Access Rules Page, page 16-10
- Looking Up Events for a Security Manager Policy, page 69-55
- Viewing Events for an IPS Signature, page 69-57
- Viewing Events for HPM Devices and Site-to-Site VPNs, page 69-58

Viewing Events for an IPS Signature

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

When an IPS device detects and reports a network intrusion by comparing incoming traffic to a configured signature, a syslog message is generated on the device. If the device is monitored by Security Manager, an incident is generated in Event Viewer after the log associated with the signature is obtained from the device. Looking up the events associated with a specific signature lets you quickly identify attacks and tune your device configuration to minimize or prevent intrusions.

To view reported network intrusion events in Event Viewer, you can select one or more entries in the Signatures policy for a device in Security Manager and navigate to the Event Viewer to view real-time and historical events.

Related Topics

- Looking Up Events for a Security Manager Policy, page 69-55
- Viewing Events for an Access Rule, page 69-56
- Viewing Events for HPM Devices and Site-to-Site VPNs, page 69-58

Step 1 (Device view) With an IPS device selected, select **IPS > Signatures > Signatures** to display the Signatures Page, page 39-4.

Step 2 Right-click the desired entry in the signatures table, or select multiple entries before right-clicking one of them, and choose one of the following commands from the **Show Events** menu:

- **Realtime**—To view real-time query results in Event Viewer for events matching this signature. Use this option to view raw events as they stream to Event Viewer.

  You can change the query criteria in the Event Monitoring window at any time, applying new parameters to alter the real-time results.

- **Historical**—To view historical query results in Event Viewer for events matching this signature.

  You can change the query criteria in the Event Monitoring window at any time, applying new parameters to alter the results.

Tips:

- If a signature is disabled, you are warned and asked if you want to proceed to event lookup.
- Events of type Packet Data and Context Data are not displayed in the query results because these events are not triggered by signature rules.

Viewing Events for HPM Devices and Site-to-Site VPNs

From Health and Performance Monitor, you can quickly access events for a monitored device or for site-to-site VPNs that have had a tunnel up/down event.

To view events for a monitored device, select a device from the All Devices, Firewall Devices, IPS Devices, Priority Devices, or custom device-related view, and with the Summary tab selected in the device details area, click the **View Events** button. Event Viewer opens and the Event Monitoring window lists events filtered by the selected device and the time period specified by the slider bar.

To view related events for a site-to-site VPN that has had a tunnel up/down event, do one of the following:

- From the Site-to-Site Tunnels view, click on the Down notification hyperlink in the Status column.
- From the Alerts view, click on the hyperlink in the Description column for tunnel up/down alerts.

Event Viewer will show IPSec VPN Events for the device within a time range depending on the polling interval for that device. If it is a priority device, the time range will be 5 minutes before until 5 minutes after the first up/down notification was received. For non-priority devices, the time range will be +/- 10 minutes instead of 5 minutes.

Related Topics

- Chapter 71, “Health and Performance Monitoring”
- Looking Up Events for a Security Manager Policy, page 69-55
- Viewing Events for an Access Rule, page 69-56
- Viewing Events for an IPS Signature, page 69-57

Examples of Event Analysis

There are many different techniques you can use to analyze and respond to events generated by your network devices. The examples in this section can help you understand some of the things you can do with the Security Manager Event Viewer.

This section contains the following topics:

- Help Desk: User Access To a Server Is Blocked By the Firewall, page 69-59

- Monitoring and Mitigating Botnet Activity, page 69-61
- Removing False Positive IPS Events from the Event Table, page 69-66

Help Desk: User Access To a Server Is Blocked By the Firewall

In this example, the help desk gets a call from a user who cannot access a server.

There are many reasons that a user might not be able to access a server, such as:

- Problems at the server's end of the network, including server down, no network connection, or the server's firewall is actively preventing access by policy.
- Problems in the network cloud between the user and the server, such as routing problems.
- Problems in the user's network, which could include workstation problems, physical problems with a network connection (for example, broken wires), problems with the switch port or wireless access point, DNS lookup failures, and so forth.

The Security Manager Event Viewer cannot identify or resolve these problems. However, it can identify whether a firewall that you control is blocking access to the server. This can help you either to rule out the firewall as being the source of the problem, or if it is blocking access, to fix the problem or to inform the user that the server is blocked by policy.

This procedure assumes that you have first determined that access to the server is not being denied by policy and that the firewall should allow access to the server.

Step 1 Ask the user for the IP address of the workstation and server.

Step 2 Open Event Viewer, for example, by selecting **Launch > Event Viewer** in Configuration Manager.

Step 3 Double-click the **Firewall Traffic Events** view to open it. Optionally, you can use the **All Device Events** view if you also want to see if there are any IPS events related to the workstation.

Tip You can also select the **Firewall Denied Events** view to see just denial events. However, you might want to see other events related to the user's workstation.

Step 4 Ask the user to retry the server access.

Step 5 Click the **Start** button, or select **View > Start**, to refresh the event table with the latest events.

Step 6 Type the user's IP address into the **Search within Results** box. The list of events is filtered as you type, and presents events in which the search string appears in any column. In the following illustration, the event list shows all events in the past 10 minutes for the IP address 10.52.150.50.

Figure 69-5 Restricting the Events List to One IP Address

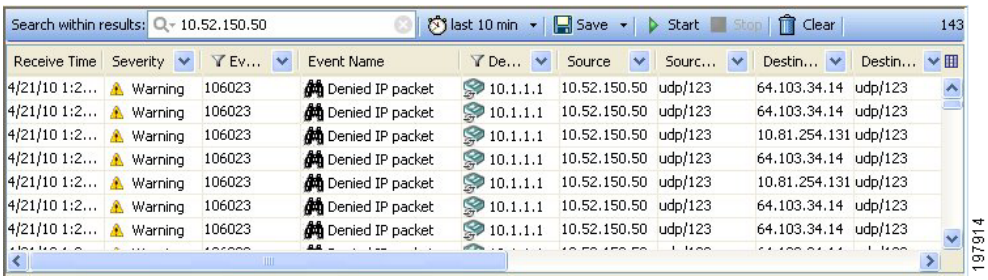

Tip You can also select the IP address from the Source column's drop-down list, and the server's IP address from the Destination column's drop-down list (or the reverse), to show only events with both the source and destination that interests you. Use the column filters if the search string does not sufficiently reduce the event list for easy analysis.

Step 7 Look for an event that indicates that traffic from the user's workstation to the server, or from the server to the workstation, was denied. Syslog **106xxx** messages indicate denial actions.

Select the event in the table and open the Event Details pane at the bottom of the window. The tabs in this pane show the complete message information and include plain-language explanations and recommended actions.

Step 8 If the event is message **106023** or **106100**, you can quickly locate the access rule that is denying the connection and fix it. You can identify whether you can look up policies from the event by looking at the Event Name cell in the table. If there is a binoculars icon before the event name, policy lookup is available. Also, if the Go To Policy command is greyed out, you cannot look up policies for that type of event.

Tip If the traffic is denied because of the implicit **deny any** rule at the end of the access list, the Go To Policy command cannot take you to the rule. For tips about rule lookup, see Looking Up a Security Manager Policy from Event Viewer, page 69-54.

a. Right-click the event and select **Go To Policy**. You are taken to Device view with the rule selected. You are notified if a matching rule cannot be found.

b. Modify the rule so that it allows the desired access. This might be as simple as deleting the rule, or you might have to add a new rule that specifically allows traffic to or from the destination server (place the permit rule above the deny rule). Your organization's security policy determines the allowable changes. For more information about configuring the access rules policy, see Configuring Access Rules, page 16-7.

c. Submit and deploy the updated configuration to the device. For more information on the deployment process, see Deploying Configurations in Non-Workflow Mode, page 8-28 or Deploying Configurations in Workflow Mode, page 8-34.

Wait for deployment to complete successfully.

Step 9 Ask the user to try to access the server again. If access is again denied, click **Start** in Event Viewer to refresh the events list and find the latest denial event.

Tip There might be more than one access rule that can deny communications with the server. The access rule policy is processed in order, top to bottom, so deleting a rule that prevents access can result in a rule that previously was not being hit suddenly becoming active. If you have a very long access rule policy, you could have several rules that you will have to remove one after the other. Alternatively, you could use the Rule Combiner tool to consolidate and simplify your access rules policy; for more information, see Combining Rules, page 12-22.

Step 10 Continue to resolve access denial events until the firewall is no longer blocking access.

Tip You can also use the Packet Tracer tool to simulate traffic going through the ASA device from the workstation to the server. In Device view, right-click the device that is denying access and select **Packet Tracer**. For more information, see Analyzing an ASA or PIX Configuration Using Packet Tracer, page 72-23.

After resolving all events, if the user still cannot reach the server, you know that the firewall is no longer one of the network elements that is blocking access. Consider other intervening network devices; perhaps a router includes an access rule that blocks the traffic.

Monitoring and Mitigating Botnet Activity

After you configure Botnet Traffic Filtering as described in Chapter 19, "Managing Firewall Botnet Traffic Filter Rules", you want to monitor it and resolve any problems identified in your network. You can use Security Manager and ASDM to monitor Botnet activity, and mitigate identified problems, as explained in the following sections:

- Understanding the Syslog Messages That Indicate Actionable Events, page 69-61
- Monitoring Botnet Using the Security Manager Event Viewer, page 69-62
- Monitoring Botnet Using the Security Manager Report Manager, page 69-64
- Monitoring Botnet Activity Using the Adaptive Security Device Manager (ASDM), page 69-64
- Mitigating Botnet Traffic, page 69-65

Understanding the Syslog Messages That Indicate Actionable Events

Botnet Traffic Filter events use syslog message numbers 338xxx. However, some messages are informational and require no action on your part.

When viewing syslogs for botnet events, you should be most concerned with the following message numbers. For information on dealing with messages that indicate blocked or allowed traffic, see Mitigating Botnet Traffic, page 69-65. For detailed descriptions of syslog messages, see the Syslog Message document for your ASA software version at http://www.cisco.com/en/US/products/ps6120/products_system_message_guides_list.html.

- **338001 to 338004**—Indicate blocked traffic that the ASA is logging, but the ASA is not stopping the traffic. These messages require immediate attention if you want to stop botnet activity that is in progress.

- **338005 to 338008**—Indicate blocked traffic that the ASA is logging and dropping. This indicates that the traffic was covered by a drop rule. Thus, your network is being protected, although you still need to disinfect the victim computer.
- **338201, 338202**—Indicate greylisted traffic that the ASA is logging but not dropping. These messages can indicate an active botnet connection that needs to be handled immediately.
- **338203, 338204**—Indicate greylisted traffic that the ASA is logging and dropping. Your network is protected from this traffic. However, if the greylisted site is legitimate, the fact that the traffic is being dropped might be a problem that requires immediate attention. You can the greylisted address to an allow list if you determine it is legitimate and redeploy the configuration, as described in Adding Entries to the Static Database, page 19-5.
- **338305 to 338307, 338310**—The ASA could not download the dynamic filter database. Ensure that you configured DNS lookup on the device, and that there is a routable network path to the Cisco Intelligence Security Operations Center. You might need to contact Cisco Technical Support.
- **338309**—The Botnet Traffic Filter license is not current, and you cannot download the dynamic database. Purchase and install the appropriate license. The Botnet Traffic Filter license is time-based, so you might have had a valid license that expired.

Monitoring Botnet Using the Security Manager Event Viewer

You can use the Event Viewer application to monitor syslog events generated by an ASA device. The Event Viewer has a predefined view that shows just botnet events.

Botnet messages are in the informational to debug severity levels and are numbered 338xxx.

Tip This procedure assumes the Event Manager service is enabled. If it is not, enable it using the **Tools > Security Manager Administration > Event Management** page.

Step 1 Open Event Viewer, for example, by selecting **Launch > Event Viewer** in Configuration Manager.

Step 2 Double-click **Botnet Events** from the list of predefined views in the left pane. You must double-click to activate the view and load it into the right pane. To verify the view has been opened, ensure that the tab name for the view in the right pane says “Botnet Events.” The following illustration shows an example of the botnet events view.

Figure 69-6 Botnet Events View in the Security Manager Event Viewer

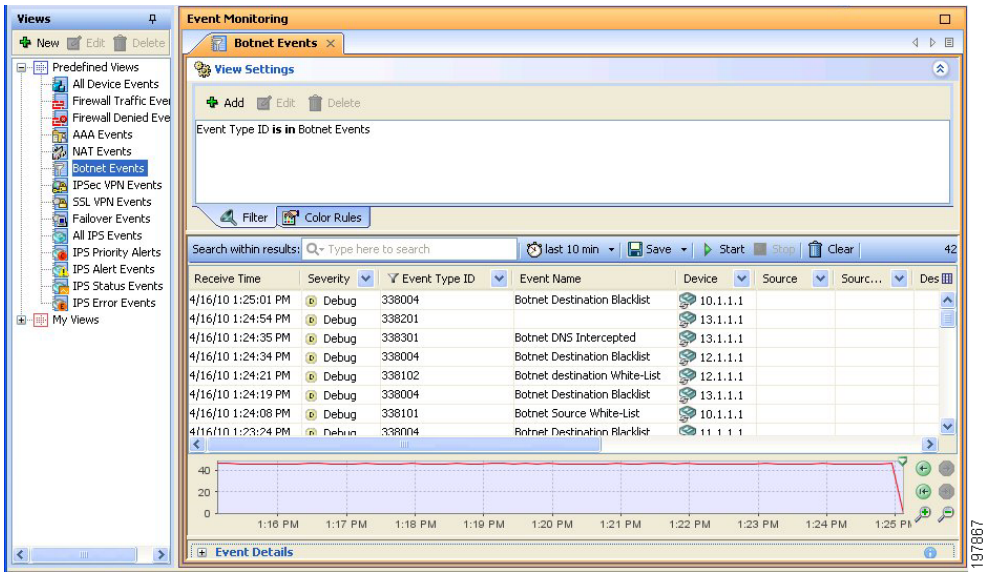

Step 3 To see the details of a specific event, select it in the table. You can then do the following:

- Double-click the event to see the tabular information presented in a more readable format.
- Open the **Event Details** section at the bottom of the window. The details pane shows information about the event organized on tabs. The Explanation and Recommended Action tabs include plain-language information about the event and what you might want to do about it.

The following illustration shows the Event Details pane for the Botnet Destination Blacklist message 338004. In this example, the recommended action is shown. The explanation for this message is "This syslog message is generated when traffic to a blocked IP address in the dynamic filter database appears." For information on dealing with this type of event, see Mitigating Botnet Traffic, page 69-65.

Figure 69-7 Botnet Event Details for Message 338004, Botnet Destination Blacklist

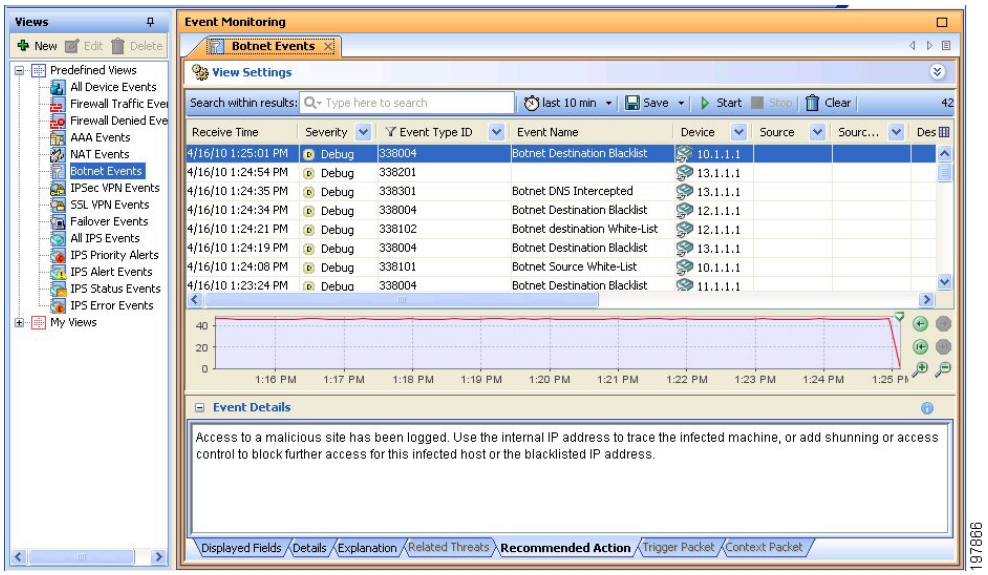

Step 4 To narrow the list of events to those generated by a single ASA, click the drop-down arrow in the Device column and select the desired device from the list. If you want to narrow the list to multiple ASAs, select Custom from the drop-down list and select the desired devices in the dialog box that appears.

You can also narrow the list using filters for any of the other columns. Filtering works the same way for all columns: either select the desired value from the drop-down list, or select Custom to create a more complex column filter.

Monitoring Botnet Using the Security Manager Report Manager

You can use the Report Manager application to generate reports on botnet activity. There are predefined reports that show the top infected hosts, the top malware ports, and the top malware sites. For a description of these reports, see Understanding Firewall Summary Botnet Reports, page 70-15.

Tip This procedure assumes the Event Manager service is enabled. If it is not, enable it using the **Tools > Security Manager Administration > Event Management** page.

Step 1 Open Report Manager, for example, by selecting **Launch > Report Manager** in Configuration Manager.

Step 2 Open the desired report from the **System > FW > Summary Botnet** folder. You can open the report by double-clicking it or by right-clicking and selecting **Open Report**.

Step 3 (Optional) Customize the report to select the desired time range and devices to include in the report. For more information, see Editing Report Settings, page 70-22.

If you want to save your custom settings to generate the report again in the future, click **Save As** to create a custom report. For more information, see Creating Custom Reports, page 70-21.

Step 4 Click **Generate Report** to retrieve the collected information and display the graphs and tabular data. For more information, see Opening and Generating Reports, page 70-20.

If you want to generate the report on a regular basis, you can configure a schedule as described in Configuring Report Schedules, page 70-33.

Monitoring Botnet Activity Using the Adaptive Security Device Manager (ASDM)

The Adaptive Security Device Manager (ASDM) includes botnet reporting features. A read-only version of ASDM is installed with the Security Manager client as a device manager, and you can start ASDM from within Security Manager.

Tip You can also install the full ASDM application separately. However, any configuration changes that you perform in ASDM are considered out-of-band changes by Security Manager and are overwritten the next time you deploy configurations from Security Manager. If you ever find a need to make configuration changes using ASDM, be sure to rediscover policies on the device in Security Manager so that Security Manager's view of the configuration is up-to-date.

Step 1 In Device view in Configuration Manager, select the ASA device.

Step 2 Select **Launch > Device Manager** to open an ASDM connection to the ASA. You are warned that you cannot make configuration changes. Click **Yes** to continue.

Step 3 In ASDM, view Botnet Traffic Filter monitoring information in the following areas:

- **Home > Firewall Dashboard** includes a Botnet Traffic Filter summary.
- **Monitoring > Botnet Traffic Filter > Reports** includes charts on the top botnet sites, ports, and infected hosts.
- **Monitoring > Logging > Log Buffer** shows historical syslog messages.
- **Monitoring > Logging > Real-Time Log Viewer** shows syslog messages as they are generated.

Tip You can also search the dynamic database on the **Configure > Botnet Traffic Filter > Botnet Database** page. This page also allows you to manually start a database download or to purge the dynamic database. These actions do not change the device's configuration and do not require policy rediscovery in Security Manager.

Mitigating Botnet Traffic

Botnet traffic mitigation is a two step process:

1. Stop traffic from your network to the botnet control site.
2. Disinfect the victim computers.

The following procedure explains the process in more detail.

Step 1 You see syslog events that indicate that packets are traveling to or from an objectionable address, typically message numbers 338001-338008 or 338201-3382004. For detailed information about these messages, see Understanding the Syslog Messages That Indicate Actionable Events, page 69-61.

Tip Messages 338201-3382004 are for greylisted traffic. You might want to first determine if the greylisted traffic is truly objectionable before stopping the traffic.

Step 2 Stop the botnet traffic:

- Messages 338005-338008 and 338203-338204 indicate that the ASA is already dropping the traffic for you. Traffic classification drop rules cover the blocked or greylisted addresses. See Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6.
- Messages 338001-338004 and 338201-338202 indicate that the ASA is logging the event but not dropping the traffic. The first order of business is to stop this traffic.

You have these options for stopping the botnet traffic if the ASA is not already dropping it because of a drop rule:

- (Preferred method.) Configure a drop rule for the botnet site and redeploy the configuration. See Enabling Traffic Classification and Actions for the Botnet Traffic Filter, page 19-6.
- (Second best method.) Log into the ASA using an SSH client, enter privileged EXEC mode, and use the **shun** command to prevent traffic to or from the botnet site. You can also issue this command through ASDM in a CLI window, but you cannot do it from Security Manager. The shun command does not create a permanent rule blocking traffic.

For example, if the botnet site is 10.1.14.14, and the internal infected computer is 10.100.10.10, issue the following commands. The first command blocks all incoming traffic from the botnet command center, the second blocks traffic from the infected computer just to the botnet site.

shun 10.1.14.14

shun 10.100.10.10 10.1.14.14

- (Not recommended.) Although the shun command is preferred, you can also create a permanent rule in the interface's access control list (ACL) that denies traffic to or from the botnet site. With the device selected in Security Manager, select **Firewall > Access Rule**, and create two rules: one that denies the botnet site as the source address, with any destination address; one that denies any source address with the botnet site as the destination address. For service, select IP so that all traffic is blocked. You must deploy the configuration for the rule to take effect.

  Creating an access rule is not the preferred method because it creates a permanent rule, whereas botnet sites are transient. Using the Botnet Traffic Filter to dynamically block botnet traffic is a better fit for this type of network attack compared to traditional access rules.

Step 3 Shut down network access for the infected computer. For example, find the switch port to which the computer is attached, and shut down the port using the switch's CLI. There might also be wireless access for the computer, so completely shutting down network access might not be a simple task.

Step 4 Inform the owner of the victim computer that it is infected and dispatch IT personnel to disinfect the computer. Tools and techniques for disinfecting a computer are outside the scope of this document.

Removing False Positive IPS Events from the Event Table

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

An IPS appliance or service module (IPS device) triggers an alarm when a given packet or sequence of packets matches the characteristics of known attack profiles defined in the IPS signatures. False positives (benign triggers) occur when the IPS reports certain benign activity as malicious. Because each event requires human intervention to diagnose, spending your time analyzing false-positive events can significantly drain resources.

Due to the nature of the IPS signatures that are used to detect malicious activity, it is almost impossible to completely eliminate false positives without severely degrading the effectiveness of the IPS or severely disrupting the computing infrastructure of an organization (such as hosts and networks). Customized tuning when an IPS is deployed minimizes false positives. Periodic re-tuning is required when the computing environment changes (for example, when new systems and applications are deployed). IPS devices provide a flexible tuning capability that can minimize false positives during steady-state operations.

An example of a false-positive is a network management station that periodically builds a network discovery map by running ping sweeps. A ping sweep triggers the ICMP Network Sweep with Echo signature (signature ID 2100). Thus, ICMP Network Sweep with Echo events that have the IP address of the network management station as the source address are actually expected and desired events.

You have the following options to remove false-positive IPS events from the event table in Event Viewer:

- **Filter out events from known "clean" sources.**

By filtering out the events, you do not stop their generation, but you also do not see them in the table. Because they are still available (you can remove the filter), you can see the events if some particular network behavior requires that you examine activity from the excluded host.

There are two main drawbacks to using this technique:

- The events are still generated, adding events to the event store.
- The filter excludes all events from a host. You cannot create a complex filter that excludes a host/signature ID pair.

The procedure below shows how to filter out events from sources that you identify as clean.

- **Create event action filter rules to stop the generation of the false-positive events.**

  Event action filter rules are the easiest way to stop generating events, and are thus preferable to editing signatures or creating custom signatures, which is a more difficult task. If you exclude a host in an event action filter rule, the IPS device does not generate alarms or log records when the host triggers the event.

  Because you can target specific signatures, rather than making a blanket-exclusion of all events from a host, you can eliminate only those events that you are certain are benign. For example, the following event filter rule removes the Produce Alert action from the ICMP Network Sweep with Echo (2100) signature for the network management station 10.100.15.75. The network management host is identified as the attacker address; the action specified in an event filter rule is actually the action that is removed from the event. Note that if you create an event action override rule to add other alert-producing actions to ICMP Network Sweep with Echo events, you must also remove the override action in this rule.

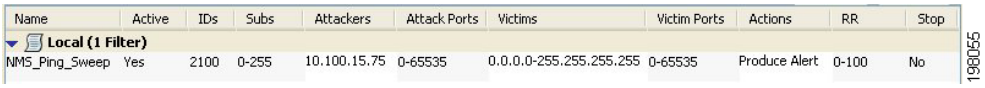

Name	Active	IDs	Subs	Attackers	Attack Ports	Victims	Victim Ports	Actions	RR	Stop
Local (1 Filter)										
NMS_Ping_Sweep	Yes	2100	0-255	10.100.15.75	0-65535	0.0.0.0-255.255.255.255	0-65535	Produce Alert	0-100	No

198055

  For more information about configuring event action filter rules, see Configuring Event Action Filters, page 40-4.

The following procedure shows how to use filtering in Event Viewer to remove false positives from the events list. It uses network/host policy objects to accomplish the filtering.

Tip By creating source or destination address filters using network/host objects, you can update the filters simply by changing the contents of the object. You do not need to add or remove filters from your views. Another advantage is that you can proactively create filters for addresses that do not currently appear in the events table; the source/destination column filter controls in Event Viewer list only those addresses that currently appear in listed events.

Step 1 Create a network/host policy object that includes the IP address of the clean hosts or networks.

a. Select **Manage > Policy Objects** to open the Policy Object Manager window (see Policy Object Manager, page 6-4).

b. Select **Networks/Hosts** from the table of contents.

c. Click the **Add Row (+)** button beneath the table of network/host policy objects, and select **Group** as the object type.

d. In the Add Network/Host Group dialog box, enter a name for the object, for example, IPS_Safe_Hosts.

e. Select **Enter IPv4 Address Information** and enter the IP address, for example, 10.100.15.75.

f. Click **Add >>** to add the IP address to the Members in Group list.

g. Click **OK** to create the object.

h. Click **Close** to close the Policy Object Manager window.

Step 2 Select **File > Submit** to submit your changes to the database (non-Workflow mode). Keep in mind that all of your configuration changes are submitted, not just the new policy object.

If you are using Workflow mode, you must submit your activity and have it approved, if necessary.

Tip Event Viewer can see only those policy objects that have been submitted to the database, so you must submit your changes before you can create a filter using the object. If you later change the object, you must also submit your changes for the filter to use the new definition of the policy object.

Step 3 Select **Launch > Event Viewer** to open the Event Viewer application.

Step 4 Create a custom view that filters out the network management station:

a. Double-click the predefined view that you want to use as the basis of your custom view, for example, **All IPS Events**. Double-clicking the view in the Views list opens the view. If you already have a custom view that you want to update, open it.

b. Click the down arrow button in the title of the Source column in the events table and select **Custom** to open the Custom Filter for Source dialog box.

Tip: You can also get to this dialog box through the View Settings pane by clicking the **Add** button, then selecting Source in the Add Custom Filter to a Column dialog box and clicking **OK**.

c. In the Custom Filter for Source dialog box, select the policy object you created and click the right-arrow button to move it to the selected list. Also, select the **Not** option next to the Condition option. The following illustration shows how the dialog box should look.

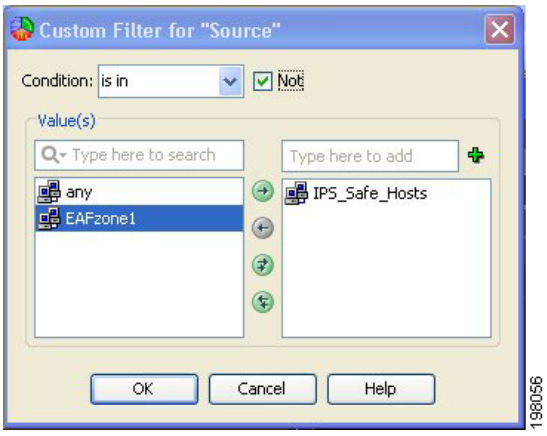

d. Click **OK**. The filter is added to the view settings and is used to remove events from the table.

e. Select **File > Save As** to save the changes as a new custom view. You are prompted for a view name and description; enter the information and click **OK**.

The following illustration shows what the view settings would look like if you started with the All IPS Events predefined view and named your new view Filtered IPS Events.

Managing Reports

Use the Report Manager application to view security and usage reports for devices and remote access IPsec and SSL VPNs.

This chapter contains the following topics:

- Understanding Report Management, page 70-1
- Overview of Report Manager, page 70-6
- Understanding the Predefined System Reports in Report Manager, page 70-13
- Working with Reports in Report Manager, page 70-19
- Scheduling Reports, page 70-32
- Troubleshooting Report Manager, page 70-36

Understanding Report Management

Use the Report Manager application to view security and usage reports for devices and remote access IPsec and SSL VPNs. These reports can provide useful information about your network.

The Report Manager aggregates information that is collected from monitored devices by the Event Manager service. Thus, to view reports about a device, you must also be monitoring the device in Event Viewer. Some statistics, such as VPN statistics, are obtained directly from the device through regular polling at five minute intervals. Aggregated data is kept for 90 days with data aggregated at 15-minute, hourly, daily, and monthly intervals; 15-minute aggregated data is kept for up to three days, hourly data up to one week. For more information about data aggregation, see Understanding Report Manager Data Aggregation, page 70-4.

You can use Report Manager to develop reports on the following:

- Adaptive Security Appliances (ASA) running ASA Software releases 8.0 and higher. ASA Software 7.x releases are also supported for VPN reports.

  Note VPN reports are not available for Cisco Catalyst 6500 Series ASA Services Modules (ASA-SM), which do not support any VPN configuration. Other types of reports are available for the ASA-SM.

- IPS devices running IPS software release 6.1 and higher (but not IOS IPS devices). This includes dedicated IPS modules installed in ASAs, routers, and switches.
- Remote access IPsec and SSL VPNs hosted on a supported ASA device.

Note Report Manager does not report on FWSM events even though Event Viewer works with FWSM.

The following topics explain Report Manager and its available reports in more detail, and also describe the other types of reports available in Security Manager:

- Understanding the Types of Reports Available in Security Manager, page 70-2
- Preparing Devices for Report Manager Reporting, page 70-3
- Understanding Report Manager Data Aggregation, page 70-4
- Understanding Report Manager Access Control, page 70-5
- Understanding the Predefined System Reports in Report Manager, page 70-13

Understanding the Types of Reports Available in Security Manager

Security Manager provides a variety of reporting capabilities. The following are the types of reports available:

- **Security and Usage Reports (Report Manager application)**—You can use the Report Manager application to view aggregated information collected by the Event Manager service from monitored devices. Some information is also obtained directly from the devices. These reports can provide information on network security and remote access IPsec and SSL VPN usage.
- **Activity (Configuration Session) Change Reports**—These reports provide detailed information about the policies changed within a specific activity (in Workflow mode) or configuration session (in non-Workflow mode). For more information, see Viewing Change Reports, page 4-16.
- **Out of Band Change Report**—These reports identify inconsistencies between the configuration that exists on a device and the configuration for the device maintained in Security Manager. You can use this information to proactively address these inconsistencies before deploying configurations, where the change will either be overwritten or the deployment will fail, depending on the behavior you select in the deployment job. For more information, see Detecting and Analyzing Out of Band Changes, page 8-45.
- **Audit Report**—This report provides information about changes to Security Manager and the objects contained in the database. The report includes information about the runtime environment, such as logins and authentication failures, changes to objects, such as activity changes and deployments, and changes to managed devices, such as inventory additions and deletions. For more information, see Generating the Audit Report, page 10-20.
- **Inventory Status**—This report provides information on policy deployment status. For more information, see Viewing Inventory Status, page 72-12.
- **Policy Discovery Status reports**—When you discover policies from a device, either while adding it to the inventory or when rediscovering policies on a managed device, the information about the policy discovery is maintained so that you can view it at a later time. For more information, see Viewing Policy Discovery Task Status, page 5-22.
- **Deployment Status reports**—When you deploy configurations to managed devices, information about the deployment is maintained so that you can view it at a later time. For more information, see Viewing Deployment Status and History for Jobs and Schedules, page 8-26.
- **Deployment and Discovery Status reports for troubleshooting**—You can export deployment and policy discovery status reports in a form suitable for sending to Cisco Technical Support (TAC) to help troubleshoot problems. You might also find these reports useful for your own purposes. For more information, see Generating Deployment or Discovery Status Reports, page 10-30.

- **Extranet VPN Configuration Summaries**—You can print, or generate a PDF file of, a summary of the configuration of an Extranet VPN. This summary can include the preshared key used for the connection. You can use this information to maintain a current record of connections between your network and the networks of partners or service providers. For more information, see Viewing a Summary of a VPN Topology's Configuration, page 25-63.
- **Policy Object Usage report**—This report shows you where a policy object is used, including instances where it is referred to by a policy or another policy object. You can use this information to help determine whether a proposed change to the object will provide the desired effect in all cases where it is used. The information is also helpful if you want to delete an object, because you cannot delete an object that is actively being used by a policy or another policy object. For more information, see Generating Object Usage Reports, page 6-15.
- **Policy Object Override report**—This report shows you all of the device-level overrides currently defined for a policy object, if the object is defined so that overrides are allowed. You can also create and delete overrides from this report. For more information, see Creating or Editing Object Overrides for Multiple Devices At A Time, page 6-19 and Policy Object Overrides Window, page 6-20.
- **Device Manager reports**—Security Manager includes read-only versions of individual device managers, such as the Adaptive Security Device Manager (ASDM), for most supported devices. You can start these device managers directly from Security Manager's Configuration Manager application and use any type of report available in those device managers. These reports are for a single device, and can augment the reports available through Report Manager. They can also provide status information for devices that are not directly supported by Event Viewer or Report Manager. For more information, see Starting Device Managers, page 72-14.

Preparing Devices for Report Manager Reporting

Before you can view reports about a device in Report Manager, you must configure the device to send events to Security Manager and configure Security Manager to monitor the device. Report Manager can provide reports only for devices you are monitoring in Event Viewer, so the device configuration for reporting is identical to the configuration for event monitoring.

Step 1 Configure the devices to send events to Security Manager. You can use Report Manager with the following types of devices:

- ASA 8.0 and higher—For the detailed configuration steps, see Configuring ASA and FWSM Devices for Event Management, page 69-28.
- IPS 6.1 and higher—For the detailed configuration steps, see Configuring IPS Devices for Event Management, page 69-29.

Step 2 Ensure that the devices are selected for event management as described in Selecting Devices to Monitor, page 69-34.

Step 3 Ensure that the Event Manager service is enabled as described in Starting, Stopping, and Configuring the Event Manager Service, page 69-30.

Understanding Report Manager Data Aggregation

Report Manager aggregates information that is collected from monitored devices by the Event Manager service. Thus, to view reports about a device, you must also be monitoring the device in Event Viewer.

Report Manager collects data using two techniques. First, the Event Manager service provides relevant events to Report Manager and then Report Manager decides if it should store those events based on the predefined reports and custom reports that are currently configured. Second, some statistics, such as VPN statistics, are obtained directly from the device through regular polling at five minute intervals.

Table 70-1 Report Manager Data Sources

Reports	Data Sources
FW Reports	
Top Sources Top Destinations Top Services	Built Syslogs: 302013,302015,302017,302020 Deny syslogs: 106001,106006,106007,106010,106011,106014,106015,106016,106017
Top Malware Sites Top Malware Ports Top Infected Hosts	BOTNET Syslogs: 338001,338002,338003,338004,338005,338006,338007,338008,338201,338202,338203,338204
IPS Reports	
All IPS Reports	All IPS Alerts
VPN Reports	

Top Bandwidth Users (Full-Client)
Top Duration Users (Full-Client)
Top Throughput Users (Full-Client)
	For ASA version 8.3 and earlier:
	`show vpn-sessiondb full svc`
	For ASA version 8.4.1 and later:
	`show vpn-sessiondb full anyconnect`

Top Bandwidth Users (IPSec-RA)
Top Duration Users (IPSec-RA)
Top Throughput Users (IPSec-RA)
	For ASA version 8.3 and earlier:
	`show vpn-sessiondb full remote`
	For ASA version 8.4.1 and later:
	`show vpn-sessiondb full ra-ikev1-ipsec`

Top Bandwidth Users (Clientless) Top Duration Users (Clientless) Top Throughput Users (Clientless)	For all ASA versions: `show vpn-sessiondb full webvpn`
User Report	All above show commands.
VPN Device Usage Report	All above show commands.

Report Manager aggregates this collected information at 15-minute, hourly, daily, and monthly intervals. Fifteen-minute aggregated data is kept a day, hourly data up to five days, and the other data for 90 days.

The aggregation schedule occurs at fixed times: 15-minute aggregation occurs at 00, 15, 30, and 45 minutes past the hour; hourly aggregation occurs on the hour (00 minutes); daily aggregation occurs at the change of day (when midnight is reached, the day is aggregated); monthly aggregation occurs at the change of the month.

The aggregation cycle has implications in what you will see in reports:

- Report data does not cover the immediate past. Instead, it covers the most recently completed whole time period of the selected duration. For example, a one-day report covers yesterday, it does not include data for today. In other words, a one day report is not the previous 24 hours starting from the time the report is generated.
- When configuring reports with a custom time period, you cannot select a time period of less than 15 minutes long. A report will always contain at least 15 minutes of aggregated data. Minute entries are rounded to the nearest aggregation time (that is, 00, 15, 30, or 45). You can configure minutes only for custom reports that start and end on the current day.

  Also, because hourly data is kept only up to five days, you can specify hours in a custom time period only for the past five days.
- You cannot generate a report for periods that are longer than the device has been monitored. For example, when you start the Event Manager service for the first time, you will not be able to generate a monthly report until after the month changes. This might be only a few days (for example, if you start the service on the twenty-ninth day of the month), or it might be almost a full month (for example, if you start the service on the first day of the month).

  The exception to this rule is the custom time period report. Custom time period reports are generated using daily aggregation data, so you can select any custom time period.

Note Be aware that your first month of aggregated data might be significantly less than one month's worth of data. If you compare monthly reports, this might appear as a significant discrepancy when in reality you are comparing 30 days of data to 15 days (as an example).

You can configure the default time interval for predefined system reports and configure time intervals in individual reports. The following topics explain the time controls:

- Configuring Default Settings for Reports, page 70-29
- Editing Report Settings, page 70-22

Understanding Report Manager Access Control

The user privileges assigned to your username control what you can do in Report Manager. If you use local users, or other types of non-ACS access control, then all users have access to Report Manager and all reports. However, the following access limits are imposed:

- You must have system administrator or network administrator privileges to configure default settings for the predefined system reports. See Configuring Default Settings for Reports, page 70-29.
- You must have system administrator or network administrator privileges to do the following to another user's schedules: see them, enable or disable them, view results generated from them, or delete them. See the following topics:
  - Viewing Report Schedules, page 70-33

  - Viewing Scheduled Report Results, page 70-35
  - Enabling and Disabling Report Schedules, page 70-35
  - Deleting Report Schedules, page 70-36
- You must have system administrator or network administrator privileges to see a list of all custom reports configured on the server and to delete another user's custom report. See Managing Custom Reports, page 70-32.

If you use ACS to control access to Security Manager, you can also control user access to Report Manager. When using ACS:

- You can control access to the Report Manager application using the View Report Manager privilege. Using this privilege, you could prevent certain users from accessing Report Manager, or create roles that allow access to Report Manager without allowing access to Event Viewer.
- Users can view reports on devices only if they have at least View privileges to the device.

For information on integrating Security Manager with Cisco Secure ACS, see the *Installation Guide for Cisco Security Manager*.

Overview of Report Manager

Use Report Manager to create security and usage reports for ASA and IPS devices, and for remote access IPsec and SSL VPNs hosted on ASA devices. For more information about supported devices and the reports you can generate using Report Manager, see Understanding Report Management, page 70-1.

To open Report Manager, do any of the following:

- Select **Start > All Programs > Cisco Security Manager Client > Report Manager** from the Windows Start menu (your exact command path might differ), or double-click the Report Manager icon on the desktop. You are prompted to log in. For more information about starting a Security Manager client application, see Logging In to and Exiting the Security Manager Client, page 1-12.
- Select **Launch > Report Manager** from the Configuration Manager or Event Viewer applications. Report Manager is opened using the same user account that you used to log into the other application.

The following illustration and subsequent list explain the basics of Report Manager.

Figure 70-1 Report Manager Main Window

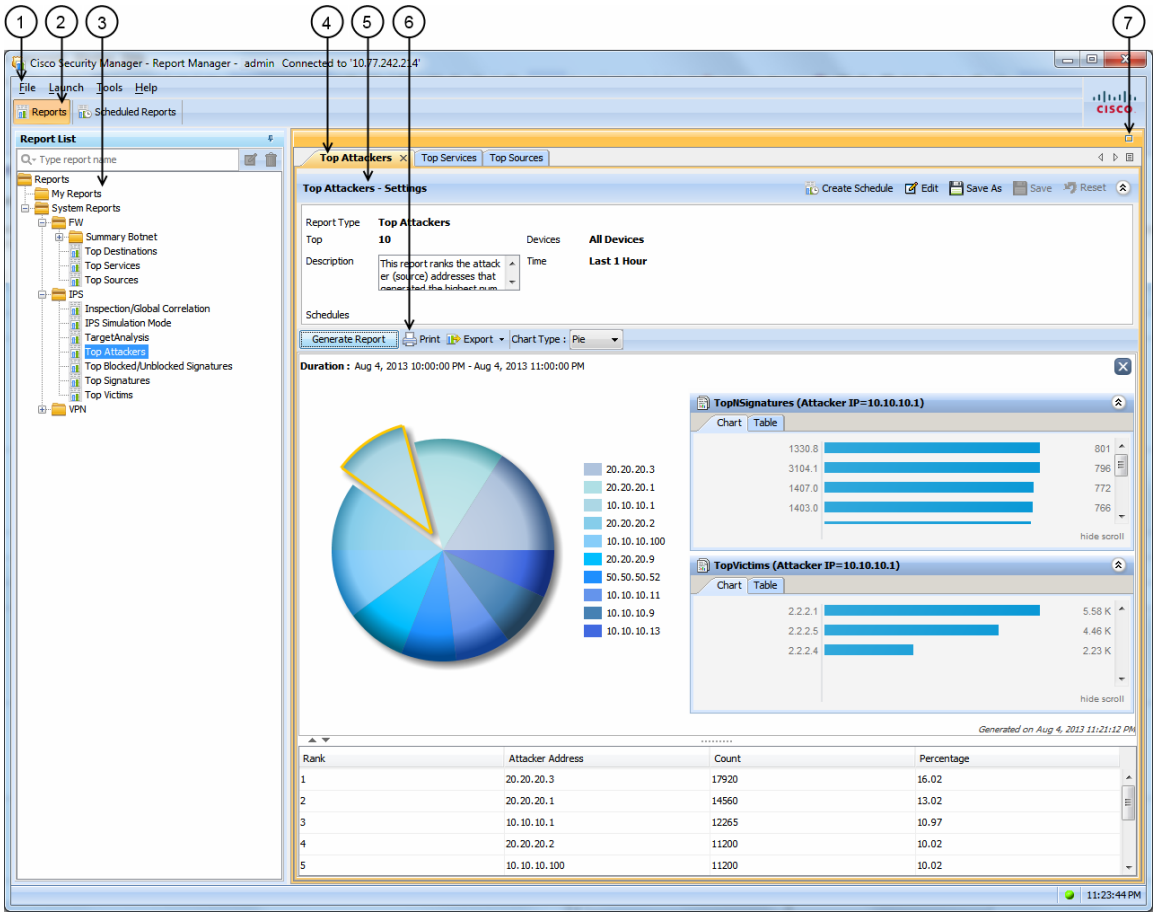

The following list explains the main Report Manager window and its call-outs in more detail.

- **Menu Bar (1)**—General commands for performing actions in Report Manager. For a description of the commands, see Report Manager Menus, page 70-8.
- **Main Window Tabs (2)**—The main window area consists of the following tabs:
  - Reports—Use the Reports tab to generate reports on demand, to create custom reports, and to perform other report-oriented tasks. The illustration above, and most of the information in this topic, relates to the Reports tab. For information on the tasks you can perform from the Reports tab, see Working with Reports in Report Manager, page 70-19.
  - Scheduled Reports—Use the Scheduled Reports tab to view and manage report schedules. For more information on the Scheduled Reports tab, see Viewing Report Schedules, page 70-33. For information on the tasks you can perform from the Scheduled Reports tab, see Scheduling Reports, page 70-32.
- **Report List (3)**—The left pane of the Reports tab is a list of reports. The list is organized into folders; the System Reports are predefined reports, whereas the My Reports folder contains the custom reports that you create. Double-click a report to open it, select the report and select **File > Open**, or right-click the report and select **Open Report**. For more information about using the report list, see Understanding the Report List in Report Manager, page 70-9.

- **Report Pane (4, 5, 6, 7)**—The right pane of the Reports tab shows the open reports. Each open report is represented on separate tabs (you can have up to five open reports). Note that you can arrange reports horizontally or vertically in this space, or even make a report float to a separate window. For more information about how you can arrange or float reports, see Arranging Report Windows, page 70-30.

  You can use the Maximize control (7) above the pane to make it take over the entire workspace (hiding the report list). After maximizing the pane, the control changes to a Restore control to return the main window to a two-pane view.

  You can use the right and left arrows, and the Show List icon button, to scroll through the open reports or to go directly to a report. However, clicking the tab with the desired report name is the easiest way to go to a report.

  The Report Pane includes these areas for each open report:

  - Report Settings pane (5)—The top part of the report shows the report settings, which are the criteria used to generate the report. You can open and close the settings pane by clicking on the heading, or on the expand/collapse icon button. The heading includes a toolbar that has commands that you can perform on the report. For more information about the settings pane, see Understanding the Report Settings Pane, page 70-10.
  - Generated Report Pane and Report Toolbar (6)—Below the settings pane is an additional toolbar used to generate and manipulate report data. Use these controls to generate the report using the criteria defined in the report settings, to print the report or to export it to PDF or CSV format, or to change the type of graphic displayed in the report.

    The bottom part of the report pane is the actual report. This area is empty until you click the Generate Report button. The top of the report shows a graphical representation of the information, the lower page shows the tabular data. For more information, see Understanding the Generated Report Pane and Toolbar, page 70-12 and Opening and Generating Reports, page 70-20.

Report Manager Menus

The following table describes the commands on the menus in Report Manager.

Table 70-2 Report Manager Menu Reference

Menu	Command	Description
File	Open	Opens the report selected in the report list on the Reports tab. See Opening and Generating Reports, page 70-20.
	Save	Saves changes made to the report settings. This command is available for custom reports only. See Saving Reports, page 70-30.
	Save As	Saves the report as a new report. Use this command to create new reports from existing reports. See Saving Reports, page 70-30.
	Close Report Close All Reports	Closes the active open report, or closes all open reports. See Closing Report Windows, page 70-31.
	Exit	Exits Report Manager.

Table 70-2 Report Manager Menu Reference (continued)

Menu	Command	Description
Launch	Dashboard Configuration Manager Event Viewer Health and Performance Monitor Image Manager	Opens the indicated Security Manager application.
Tools	Default Report Settings	Configures the default settings for predefined system reports. See Configuring Default Settings for Reports, page 70-29.
	Custom Report List	Displays all custom reports configured on the server, not just those that you created. You can manage reports from this window. See Managing Custom Reports, page 70-32.
Help	Help about this page	Opens the online help to a topic relevant to the page currently displayed in the main window.
	About Report Manager	Displays copyright, version, and licensing information for the application.

Understanding the Report List in Report Manager

The left pane of the Reports tab in Report Manager displays a list of available reports, as shown in the following illustration.

Figure 70-2 Report Manager Report List

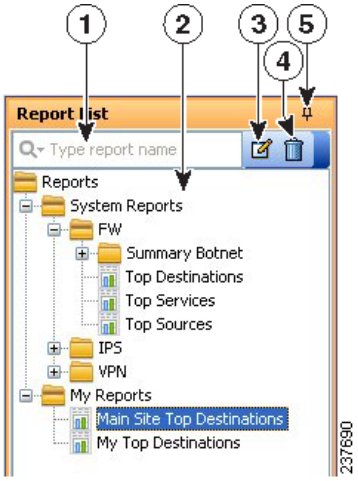

The Report List includes the following controls (illustration call-outs cited):

- **Quick Filter search box (1)**—Use the quick filter search box to search for reports in the list. The list is filtered as you type, although folders are not opened automatically. The default is to search for the text string anywhere in the report name. However, you can click the down arrow in the Quick Filter box to select a variety of options to change how your search string is evaluated.
- **List of reports (2)**—The list is organized into folders; the System Reports are predefined reports (explained in Understanding the Predefined System Reports in Report Manager, page 70-13), whereas the My Reports folder contains the custom reports that you create. Double-click a report to open it, or select the report and select **File > Open**. For more information, see Opening and Generating Reports, page 70-20.
- **Right-click shortcut menu (not shown)**—If you right-click on a report, you get a list of additional commands that you can perform, such as opening the report, creating a schedule, or saving the report as a new report.
- **Edit button (3)**—Click the Edit button to change the name of the selected custom report. You can edit custom reports only. For more information, see Renaming Reports, page 70-31.
- **Delete button (4)**—Click the Delete button to delete the selected custom report. You can delete custom reports only. For more information, see Deleting Reports, page 70-32
- **Push Pin button (5)**—Click the Push Pin icon to control whether the report list pane is opened or closed. If the pin is vertical, the report list remains open unless you maximize the report pane (the right pane). If the pin is horizontal, the report list collapses to the left margin, and you must click the Report List heading in the left margin to open the list.

Related Topics

- Overview of Report Manager, page 70-6
- Understanding Report Management, page 70-1
- Working with Reports in Report Manager, page 70-19
- Viewing Report Schedules, page 70-33
- Scheduling Reports, page 70-32
- Arranging Report Windows, page 70-30

Understanding the Report Settings Pane

The top part of the right side of the Reports tab, with a report open, shows the report settings. These settings define the criteria that are used to generate a report. The following illustration shows an example of the report settings pane.

Figure 70-3 Report Manager Report Settings

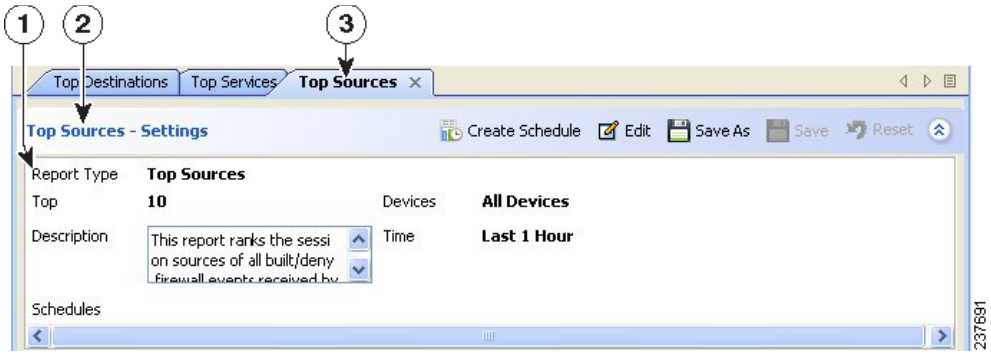

The Report List includes the following controls (illustration call-outs cited):

- **Report tab (3)**—Although not part of the settings per se, each report appears on its own tab. The settings are the top part of the tab. If you right-click the tab itself, you get a menu of commands that allow you to arrange report windows. For more information, see Arranging Report Windows, page 70-30.
- **Heading and toolbar (2)**—The top of the settings pane includes the heading (for example, Top Sources - Settings) and a row of buttons for manipulating the settings. You can open or close the pane by clicking the heading or the up arrow button in the far right of the toolbar. The other buttons have the following functions:
  - Create Schedule button—Creates a new schedule for automatically generating reports based on these settings. For more information, see Configuring Report Schedules, page 70-33.
  - Edit button—Edits the report settings. For more information, see Editing Report Settings, page 70-22.
  - Save As button—Saves the report as a new report. If you edit the settings for a predefined system report, and you want to save your changes, you must use Save As to create a custom report. For more information, see Saving Reports, page 70-30 and Creating Custom Reports, page 70-21.
  - Save button—Saves changes to the settings. You can save changes for custom reports only. For more information, see Saving Reports, page 70-30.
  - Reset button—Resets the settings to the last saved values.
  - Expand/Collapse button (double up/down arrows)—Toggles between opening and closing the report setting pane.
- **Settings display (1)**—Below the heading and toolbar is a summarization of the report settings. Information includes the type of report, the devices included in the report, the time range, a description, the schedules defined for the report, and other properties unique to the report.

  To change the description, type your changes directly into the Description edit box.

Related Topics

- Opening and Generating Reports, page 70-20
- Understanding the Generated Report Pane and Toolbar, page 70-12
- Overview of Report Manager, page 70-6
- Understanding Report Management, page 70-1
- Working with Reports in Report Manager, page 70-19

- Viewing Report Schedules, page 70-33
- Scheduling Reports, page 70-32

Understanding the Generated Report Pane and Toolbar

The bottom part of the right side of the Reports tab, with a report open, shows the generated report and report toolbar. This pane displays the results of clicking the Generate Report button.

The following illustration shows an example of the generated report pane and its associated report toolbar.

Figure 70-4 Report Manager Generated Report Pane and Toolbar

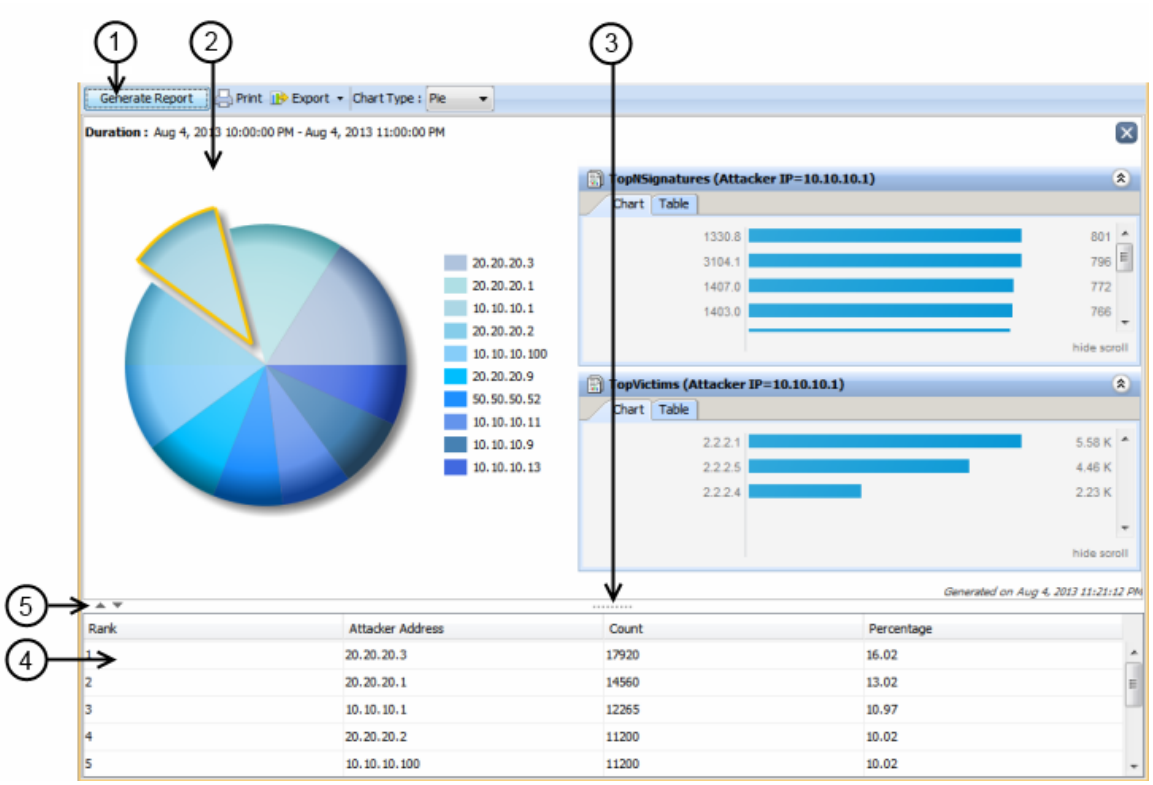

The Report List includes the following controls (illustration call-outs cited):

- **Report toolbar (1)**—The top of the generated report pane is a row of controls for generating and manipulating reports. The controls have the following functions:
  - Generate Report button—Generates a report based on the criteria defined in the report settings (in the upper pane). For more information, see Opening and Generating Reports, page 70-20.
  - Print button—Prints the generated report. For more information, see Printing Reports, page 70-27.
  - Export button—Exports the report. Click the down arrow in the button and select the type of file you want to create: **As PDF** (for Adobe Acrobat) or **As CSV** (for comma-separated values). For more information, see Exporting Reports, page 70-28.

  - Chart Type—Determines the type of graphical chart displayed in the upper part of the report, typically pie, bar, and XY (linear) graphs are available. In some cases, you do not have a choice of chart types. For more information, see Opening and Generating Reports, page 70-20.

- **Graphical view (2, 3, 5)**—The top part of the generated report shows a graphical, color-coded view of the report data, and includes a legend that explains the colors. Also included is the date and time the report was generated.

  Note For the Top Destinations, Top Services, and Top Sources firewall reports and the Top Attackers, Top Signatures, and Top Victims IPS reports, you can click on a data point in the Pie, XY, or Bar graph to see additional details for that data point. For example, if you click on a slice in the Pie chart for the Top Signatures report, you will see report details for the top attackers and top victims for the selected signature. For more information, see Drilling Down into Report Data, page 70-26.

  At the bottom of the graphical view are the following controls:

  - Up and Down arrows (5)—These icon buttons, to the left of the graphic, allow you to open and close the graphical part of the report.
  - Window size control (3)—If you hover the mouse pointer over the vertical dashes below the graphic in the center of the window, you can click and move the pointer to change the size of the graphical portion of the report. The graphic is automatically resized as you increase or decrease the size of the area. In fact, you can hover over any part of the top of the table to access this control.

- **Tabular view (4)**—The bottom part of the report is a table that presents the data collected for the report and used to produce the graphic. The columns in the table vary according to the type of report.

  You can click a heading to sort the table by the column. There are three sort orders, and clicking the column heading cycles through these orders with an arrow indicating the sort order: ascending (up arrow), descending (down arrow), and no sort (empty). You can use Ctrl+click to create a second sort order on a separate column, which has an effect only if the first sort column repeats one or more entry. Numbers indicate whether the column is the first, second, third, and so on, sort criteria.

Related Topics

- Understanding the Report Settings Pane, page 70-10
- Overview of Report Manager, page 70-6
- Understanding Report Management, page 70-1
- Working with Reports in Report Manager, page 70-19
- Viewing Report Schedules, page 70-33
- Scheduling Reports, page 70-32

Understanding the Predefined System Reports in Report Manager

Report Manager includes several predefined system reports that you can use to analyze your network. You can customize these reports to focus on specific sets of devices and time periods or to focus on other configurable parameters.

This section contains the following topics:

- Understanding Firewall Traffic Reports, page 70-14
- Understanding Firewall Summary Botnet Reports, page 70-15
- Understanding VPN Top Reports, page 70-16
- Understanding General VPN Reports, page 70-16
- Understanding IPS Top Reports, page 70-17
- Understanding General IPS Reports, page 70-19

Understanding Firewall Traffic Reports

Report Manager includes predefined system reports that you can use to identify the top destinations, services, and sources for firewall ACL events. The statistics are based on the events collected by the Event Manager service (as displayed in Event Viewer).

The following reports are available in the **System Reports > FW** folder.

- **Top Destinations**—This report ranks the session destinations of all built/deny firewall events received by Security Manager. The report shows the destination IP address, the count of the number of events for each address, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific destination to see report information about the top sources and top services associated with that destination (see Drilling Down into Report Data, page 70-26).
- **Top Sources**—This report ranks the session sources of all built/deny firewall events received by Security Manager. The report shows the source IP address, the count of the number of events for each address, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific source to see report information about the top destinations and top services associated with that source (see Drilling Down into Report Data, page 70-26).
- **Top Services**—This report ranks the destination services of all built/deny firewall events received by Security Manager. TCP and UDP services include the port number. The report shows the service, the count of the number of events for each service, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific service to see report information about the top destinations and top sources associated with that service (see Drilling Down into Report Data, page 70-26).

The parameters used to define the number of addresses or services to included in the report and the reporting time period are defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29.

You can also edit the report settings and create custom versions of the reports. You can narrow the reports to focus on specific sets of source or destination addresses or services, or on just permit or deny actions, or limit the report to focus on a sub-set of firewall devices, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Understanding Firewall Summary Botnet Reports

Report Manager includes predefined system reports that you can use to analyze botnet traffic filtering. The statistics are based on the botnet events collected by the Event Manager service (as displayed in Event Viewer) for blocked and gray-listed sites.

For more information about botnet, see Chapter 19, "Managing Firewall Botnet Traffic Filter Rules".

The following reports are available in the **System Reports > FW > Summary Botnet** folder.

- **Top Infected Hosts**—This report ranks the top infected hosts for traffic originating from infected hosts to blocked or gray-listed sites based on all botnet events received by Security Manager. The report shows the IP address of the infected host with the firewall interface name on which the event was detected in parentheses, the count of the number of connections logged to blocked or gray-listed sites for each address, the count of the number of connections that were blocked (dropped) by botnet traffic filtering, and the percentage of the count compared to the sum of all counts in the report.
- **Top Malware Ports**—This report ranks the top destination ports for traffic originating from infected hosts to blocked or gray-listed sites based on all botnet events received by Security Manager. The report shows the destination malware port, the count of the number of connections logged to blocked or gray-listed sites for each port, the count of the number of connections that were blocked (dropped) by botnet traffic filtering, and the percentage of the count compared to the sum of all counts in the report.
- **Top Malware Sites**—This report ranks the top botnet sites (blocked or gray-listed sites) for all inbound and outbound sessions based on all botnet events received by Security Manager. The report shows the following information:
  - IP Address—The IP address that is indicated as the malicious host in botnet events, either on the block list or the grey list.
  - Malware Site—The domain name or IP address in the dynamic filter database to which the traffic was initiated.
  - List Type—Whether the site is on the block list or the grey list.
  - Connections Logged—The count of the number of connections logged or monitored for each site.
  - Connections Blocked—The count of the number of connections that were blocked (dropped) by botnet traffic filtering for each site.
  - Threat Level—The botnet threat level for the site, from very low to very high, or none.
  - Category—The category of threat the site poses as defined in the botnet database, such as botnet, Trojan, spyware, and so on.

The parameters used to define the number of hosts, ports, or sites in the report and the reporting time period are defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29. You can also edit the report settings and create custom versions of the reports, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Understanding VPN Top Reports

Report Manager includes predefined system reports that you can use to identify the top remote access VPN users based on bandwidth usage, duration of connection to your network, and data throughput. Separate reports are provided based on the type of connection made by the user.

These reports are available in the **System Reports > VPN** folder in the **AnyConnect (SSL) Remote Access VPN**, **Cisco VPN Client (IPsec) Remote Access VPN**, and **Clientless SSL VPN**.

The following reports are available in each folder. Each report is specific to the connection type indicated by the folder name and also included in parentheses in the report name.

- **Top Bandwidth Users**—This report ranks the VPN users who consumed the most bandwidth. The report shows the usernames, the bandwidth in total number of bytes sent and received, and the percentage of reported bandwidth used by each user.
- **Top Duration Users**—This report ranks the VPN users who remained connected to the network for the longest time. The report shows the usernames, the connection duration time in *days hours:minutes:seconds* format, and the percentage of the reported duration by each user. The chart shows duration in seconds.
- **Top Throughput Users**—This report ranks the VPN users who sent and received data at the highest throughput rate. The report shows the usernames, the throughput for each user in kbps, and the percentage of reported throughput by each user. The throughput is calculated as 8.0*(bandwidth of the user in bytes)/(duration for which the user is connected in seconds*1000.0).

The parameters used to define the number of users included in the report and the reporting time period are defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29. You can also edit the report settings and create custom versions of the reports, including focusing on specific users, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Understanding General VPN Reports

Report Manager includes predefined system reports that you can use to analyze general remote access VPN usage in your network. These reports are not specific to the connection types used in the VPN.

The following reports are available in the **System Reports > VPN** folder.

- **Connection Profile Report**—This report provides a count of user, session, and summary of the bandwidth utilization and throughput usage for each remote access connection profile.

  The default report contains this information for all devices for the previous hour. You can customize the report in several different ways (see Editing Report Settings, page 70-22).
- **User Report**—This report provides a summary of the bandwidth utilization, connection duration and throughput usage for each remote access VPN user. The report shows the usernames, the bandwidth in total number of bytes sent and received, the connection duration time in *days hours:minutes:seconds* format, and the throughput for each user in kbps. The throughput is calculated as 8.0*(bandwidth of the user in bytes)/(duration for which the user is connected in seconds*1000.0).

  Beginning with Security Manager 4.7, the User Report provides both **user-level details** and **session-level details**:

- **User-Level Details**—For a particular user, the **user-level details** represent the combined value of all that user's sessions: Username, Total no. of Sessions, Bandwidth, Duration, and Throughput.

Note Starting from Cisco Security Manager 4.13, the Public IP and Assigned IP details are also displayed as part of the User-Level Details in the general VPN report.

- **Session-Level Details**—Expanding the tree displays the **session-level details** for each session that a particular user has a VPN connection with; the session-level details encompass the Session ID, Login Time, Logout Time, Bandwidth, Throughput, and Duration of the Session. (Here the logout time is calculated by using the formula **Logout Time = Login Time + Duration**.)

The default report includes information for all connection technologies and all users. You can customize the report to focus on a single technology type or one or more specific users (see Editing Report Settings, page 70-22).

The "Criteria" section of the User Report has filters for Technology (All, Clientless, Full Client, and IPSec RA), Username, and User Session Duration (<= and >= in hours).

- **VPN Device Usage Report**—This report provides a summary of the usage statistics for each device that hosts remote access VPN connections. The report shows the device (using the Security Manager display name), the average number of users logged into the VPN at any given time during the reported time range, the total bandwidth of all users in the VPN in bytes (sent and received), the total connection duration time in *days hours:minutes:seconds* format, and the average throughput in kbps at any given time during this report period.

  The default report includes information for all connection technologies. You can customize the report to focus on a single technology type (see Editing Report Settings, page 70-22).

The parameter used to define the reporting time period is defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29. You can also edit the report settings and create custom versions of the reports, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Understanding IPS Top Reports

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Report Manager includes predefined system reports that you can use to analyze top attackers, victims, and signatures for IPS alerts in your network.

The following reports are available in the **System Reports > IPS** folder.

- **Top Attackers**—This report ranks the attacker (source) addresses that generated the highest numbers of recorded IPS alerts. The report shows the attacker IP address, the count of the number of alerts for each address, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific attacker to see report information about the top signatures and top victims associated with that attacker (see Drilling Down into Report Data, page 70-26).

The default report includes information for all attackers, victims, and signatures for both blocked and unblocked actions. You can customize the report to focus on subsets of attackers, victims, or signatures, or limit the analysis to blocked only or unblocked only actions (see Editing Report Settings, page 70-22).

- **Top Victims**—This report ranks the victim (destination) addresses that generated the highest numbers of recorded IPS alerts. The report shows the victim address, the count of the number of alerts for each address, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific victim to see report information about the top signatures and top attackers associated with that victim (see Drilling Down into Report Data, page 70-26).

  The default report includes information for all attackers, victims, and signatures for both blocked and unblocked actions. You can customize the report to focus on subsets of attackers, victims, or signatures, or limit the analysis to blocked only or unblocked only actions (see Editing Report Settings, page 70-22).

- **Top Signatures**—This report ranks the signatures that fired the highest numbers of alerts. The report shows the signature ID number, the name of the signature, the count of the number of alerts for each signature, and the percentage of the count compared to the sum of all counts in the report. You can click on a data point in the Pie, XY, or Bar graph that represents a specific signature to see report information about the top victims and top attackers associated with that signature (see Drilling Down into Report Data, page 70-26).

  The default report includes information for all attackers, victims, and signatures for both blocked and unblocked actions. You can customize the report to focus on subsets of attackers, victims, or signatures, or limit the analysis to blocked only or unblocked only actions (see Editing Report Settings, page 70-22).

- **Top Blocked/Unblocked Signatures**—This report ranks the signatures that blocked the highest numbers of attacks. The report shows the signature ID number, the name of the signature, the count of the number of alerts for each signature, and the percentage of the count compared to the sum of all counts in the report.

  The default report shows blocked actions only. However, you can customize the report to show unblocked only or a combination of blocked and unblocked actions (see Editing Report Settings, page 70-22).

  If you want to see blocked or unblocked lists that are limited to specific attacker or victim addresses, or to a subset of signatures, use the Top Signatures report instead of the Top Blocked/Unblocked Signatures report. Customize the report to show blocked only or unblocked only signatures.

- **IPS Target Analysis**—This report provides the top targets by signature and frequency of attack. The report shows the signatures that generated the alerts, the number of alerts, and the victim IP address, and is based on an aggregated view of the Top Signatures and Top Victims reports. The report contains up to ten signatures and five attackers. The information is plotted on a scatter plot, which is the only graphical representation available for the report.

The parameters used to define the number of addresses or signatures to included in the report and the reporting time period are defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29. You can also edit the report settings and create custom versions of the reports, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Understanding General IPS Reports

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

Report Manager includes predefined system reports that you can use to analyze general IPS activity in your network.

The following reports are available in the **System Reports > IPS** folder.

- **Inspection/Global Correlation**—This report provides a comparison of alerts generated by global correlation against alerts generated by traditional IPS inspection. The report shows the number and percentage of alerts per IPS inspection method (either Global Correlation or Inspection).
- **IPS Simulation Mode**—This report provides a comparison of alerts in inline (IPS) and promiscuous (IDS or IPS simulation) modes. The report shows the number and percentage of alerts based on mode, either Non Simulation Count (inline) or Simulation Mode Count (promiscuous). The IPS sensor cannot directly block attacks that occur in promiscuous mode.

When working with IPS events, the Report Manager component of Cisco Security Manager reports events individually; the Event Viewer component of Cisco Security Manager displays alerts. In the Event Viewer component, the IPS Summarizer groups events into a single alert, thus decreasing the number of alerts that the IPS sensor sends out.

Tip Cisco IPS Manager Express (IME) and Cisco Security Manager do not summarize events in precisely the same way.

The parameter used to define the reporting time period is defined in the system defaults as described in Configuring Default Settings for Reports, page 70-29. You can also edit the report settings and create custom versions of the reports, as described in the following topics:

- Editing Report Settings, page 70-22
- Creating Custom Reports, page 70-21

Working with Reports in Report Manager

Use the Report Manager application to view security and usage reports for devices and remote access IPsec and SSL VPNs. The following topics explain the basics of creating reports. For information on working with report schedules, see Scheduling Reports, page 70-32.

This section contains the following topics:

- Opening and Generating Reports, page 70-20
- Creating Custom Reports, page 70-21
- Editing Report Settings, page 70-22
- Drilling Down into Report Data, page 70-26
- Printing Reports, page 70-27
- Exporting Reports, page 70-28
- Configuring Default Settings for Reports, page 70-29

- Arranging Report Windows, page 70-30
- Saving Reports, page 70-30
- Renaming Reports, page 70-31
- Closing Report Windows, page 70-31
- Deleting Reports, page 70-32
- Managing Custom Reports, page 70-32

Opening and Generating Reports

Reports are not static. When you open a report, it contains no data, although it does contain settings that define the data that shall be used to generate the report. Thus, to view a report, you need to open it and then generate it. This procedure explains the process.

Related Topics

- Overview of Report Manager, page 70-6
- Creating Custom Reports, page 70-21
- Arranging Report Windows, page 70-30
- Troubleshooting Report Manager, page 70-36

Step 1 In Report Manager, do one of the following to open a report:

- Double-click the name of the report in the report list (in the left pane).
- Select the report in the reports list and select **File > Open**.
- Right-click a report in the reports list and select **Open Report**.

The report opens with the report settings pane open and the report content area empty.

Tip You can have five reports at most open at one time. You can also collapse the report settings pane to provide more room for viewing the generated report by clicking on any area of the settings toolbar that is not a button that performs another function.

Step 2 (Optional) Verify that the report settings contain the desired values, for example, the desired time window for the report. The settings for system reports are based on the system defaults (which you can configure as described in Configuring Default Settings for Reports, page 70-29). The settings for custom reports are those that were last saved for the report.

If you need to change the settings, click the **Edit** button in the settings toolbar and make your changes in the Edit Settings dialog box. For more information, see Editing Report Settings, page 70-22.

Tip Be sure to save your changes if you want to make them permanent. If you change the settings for a system report and you want to preserve them, you must use Save As to create a new custom report; you cannot change the settings of a system report from the default settings.

Step 3 Click the **Generate Report** button below the settings pane to retrieve the report data from the reporting database and to display the resulting information. The information is displayed in two formats:

- Graphical—A graphical representation of the data is shown in the top part of the report. You can select different types of graphics from the **Chart** menu above the report data: pie, XY (for linear graphs), and bar. If there are more than 10 items in the report (for example, you configured a Top report to show 25 values), all values after the tenth are summarized in the chart as "others."

  Note For the Top Destinations, Top Services, and Top Sources firewall reports and the Top Attackers, Top Signatures, and Top Victims IPS reports, you can click on a data point in the Pie, XY, or Bar graph to see additional details for that data point. For example, if you click on a slice in the Pie chart for the Top Signatures report, you will see report details for the top attackers and top victims for the selected signature. For more information, see Drilling Down into Report Data, page 70-26.

  Some reports, such as the IPS Target Analysis report, use scatter plots. For these reports, you do not have the option to select a different graphic type.

- Tabular—The table below the graphic lists the data used to generate the graphic. The table has different columns based on the type of report. Following are some typical columns; for more detailed information about the content of each report, see Understanding the Predefined System Reports in Report Manager, page 70-13.
  - Rank—The order of the information by magnitude. For example, for a firewall top destinations report, a rank of 1 indicates that the destination is the most used in the evaluated events.
  - (Name of reported characteristic)—There is always a column whose name is based on the characteristic targeted by the report, for example, Source/Destination (IP addresses), Service (protocol and port), or User (usernames).
  - Count—The number of times the item appears in an event or related statistic.
  - Percentage—The ratio of the reported characteristic to the total sum of that characteristic in the report. The ratio includes only those numbers included in a report, so for example, you could get a different percentage for the same item in a top 10 verses a top 25 report.

Step 4 (Optional) If desired, you can print the report or export it to a PDF or comma-separated values (CSV) file.

- To print the report, click the **Print** button and select your printer. For more information, see Printing Reports, page 70-27.
- To export the report, click the **Export** button and select the file type, PDF or CSV. For more information, see Exporting Reports, page 70-28.

Tip The report data is not preserved when you close the report. If you want to keep the displayed information, you must print or export the report.

Creating Custom Reports

You can create custom reports to target specific characteristics that require regular analysis or presentation. For example, you might want to create separate Top Destination firewall reports for different groups of firewall devices so that you can separately analyze activity in separate physical sites. You can also use custom reports to analyze sources, destinations, or services that otherwise do not make it into the top reports.

Tip It might take up to one hour for data to be available for a newly-created custom report. If you get the message that no records are found after creating the report, wait an hour and then ensure that the time span for the report is Last 1 Hour.

Related Topics

- Opening and Generating Reports, page 70-20
- Overview of Report Manager, page 70-6

Step 1 Select the report on which you want to base your custom report in the reports list. Open it by double-clicking it, by selecting it and selecting **File > Open**, or by right-clicking and selecting **Open Report**.

Step 2 Click the **Edit** (pencil) button in the settings toolbar to open the Edit Settings dialog box.

Note Do not click the Edit button above the reports list. That edit button allows you to change the name of the report only.

The Edit Settings dialog box is divided into two panes. The left pane lists the available settings pages; the right pane shows the settings for the page selected in the left pane.

Step 3 Configure the settings so that they define the desired report parameters. For more information, see Editing Report Settings, page 70-22.

Step 4 Click the **Save As** button in the settings toolbar, or select **File > Save As**.

Step 5 Enter the name of the report and optionally a description and click **OK**.

Report names can be up to 64 characters and contain alphanumeric characters, spaces, hyphens (-), and the underscore character (_). The description can be up to 1024 characters.

Editing Report Settings

You can change the settings that define the criteria used to generate a report. For custom reports, you can save your changes.

For predefined system reports, you cannot directly save your changes. Instead, you can use Save As to create a new custom report using the updated settings. You can also change the default settings used in all predefined system reports rather than editing report settings, as described in Configuring Default Settings for Reports, page 70-29.

Related Topics

- Opening and Generating Reports, page 70-20
- Overview of Report Manager, page 70-6
- Creating Custom Reports, page 70-21
- Understanding Report Manager Data Aggregation, page 70-4
- Arranging Report Windows, page 70-30

Step 1 In Report Manager, open the report whose settings you want to change. To open the report, double-click it, select it and select **File > Open**, or right-click it and select **Open Report**.

The report opens with the settings pane open at the top of the report. The settings pane shows the type of report, the devices included in the report, the time range, a description, the schedules defined for the report, and other properties unique to the report.

Step 2 (Optional) Change the Description by typing into the Description edit box in the report settings pane.

Step 3 Click the **Edit** (pencil) button in the settings toolbar to open the Edit Settings dialog box.

Note Do not click the Edit button above the reports list. That edit button allows you to change the name of the report only.

The Edit Settings dialog box is divided into two panes. The left pane lists the available settings pages; the right pane shows the settings for the page selected in the left pane.

Step 4 Edit the settings on the desired pages as follows:

- **Devices**—To change which monitored devices are included in the report. The default is **All Devices**.

  If you want the report to reflect a subset of monitored devices, select **Filter Devices** and then select the desired devices or created contexts from the list. If a device is in italics, it means that the device is not currently selected for monitoring in Event Viewer; you can select these devices, and the report will include data for the device if it was monitored during the selected time period. You can select a folder to select all the devices in the folder.

  The device list is pre-filtered to show devices of the appropriate type only. For example, if you are editing the settings for a firewall report, IPS devices do not appear in the list of selectable devices.

  Note Starting from Cisco Security Manager 4.10, all the created contexts for ASA 9.5(2) and above will be listed under **Filter Devices**.

- **Time**—To change the time span used to select events and data to include in the report. The time is based on the Security Manager server time. You can select one of the following options to define the time span:
  - Last 1 Hour—The last full one hour on the zeros, for example, if the current time is 11:45 AM, the Last 1 Hour report shows data from 10:00 to 11:00.
  - Last 1 Day—The last full day, from midnight to midnight. For example, if the current day is Tuesday, the Last 1 Day report shows data from Monday.
  - Last 1 Week—The previous Monday through Sunday.
  - Last 1 Month—The previous month. For example, if the current date is September 29, the Last 1 Month report shows data from August.
  - Custom—Use the Start Date and End Date calendars to select the desired starting and ending times for the report. Click the down arrow, select the desired day and time, then click OK in the calendar widget. Reportable data is kept for 90 days, so you cannot select a date more than 90 days into the past. Additionally, you cannot specify a time if you select a start date more than five days into the past. If you select the current date for the start date, you can also specify minutes for both starting and ending dates, but because report data is aggregated every 15 minutes at 00, 15, 30, and 45 minutes past each hour, minute entries are rounded to the nearest of these figures. The allowed time selection is based on how data is aggregated, as explained in Understanding Report Manager Data Aggregation, page 70-4.

- **Criteria**—To change the other criteria used to define the report. The attributes available on the Criteria settings page are variable. In some cases, there are no selectable criteria. Following is a list of the possible criteria:
  - Top (All "Top" reports.)—The number of items targeted by the report to include. For example, the Top 10 firewall destinations returns the 10 most frequent destinations for firewall events in the configured time range. Select 10, 20, 25, or 50.
  - Service (Firewall reports except Botnet)—The services to include in the report. To specify services, click the Edit button next to the field and select the desired service policy objects. You can select multiple objects.
  - Source IP, Destination IP (Firewall reports except Botnet)—The source and destination IP address fields are separate, but they are functionally the same. They define the IP addresses for sources or destinations to include in the report. You can enter individual addresses, such as 10.100.10.10, or address ranges, such as 10.100.10.10-10.100.10.20. Both IPv4 and IPv6 addresses are accepted. Separate multiple addresses with commas.

    You can click the Edit button next to the field to open a dialog box where you can more easily create complex lists of addresses and address ranges. However, you cannot use network/host objects to define addresses.

    Note Do not specify values for all of the Service, Source IP, and Destination IP criteria in a single report. You can specify the criteria on which the report is based (for example, Service for the Top Services report) plus one other criteria. If you specify all three values, the report will always contain no data.

  - Permit/Deny (Firewall reports except Botnet)—The action reflected in the event, either permitting the matching traffic (Permit), denying the matching traffic (Deny), or either (All). The default is All.
  - Signature ID (IPS top attackers, top signatures, top victims)—The signatures to include in the report. To specify signatures, click the Edit button next to the field and select the desired signatures. You can select a folder to select all signatures in the folder.

    Note In the predefined system reports, you cannot specify values for all three of the Signature ID, Attacker IP, and Victim IP criteria. You can specify a value for the key attribute of the report (for example, Victim IP for the top victims report) plus one of the other values. If you want to configure values for all three criteria, you must create a custom report.

  - Attacker IP, Victim IP (IPS top attackers, top signatures, top victims)—The attacker and victim IP address fields are separate, but they are functionally the same. They define the IP addresses for attackers (sources) or victims (destinations) to include in the report. You can enter individual addresses, such as 10.100.10.10, or address ranges, such as 10.100.10.10-10.100.10.20. Both IPv4 and IPv6 addresses are accepted. Separate multiple addresses with commas.

    You can click the Edit button next to the field to open a dialog box where you can more easily create complex lists of addresses and address ranges. However, you cannot use network/host objects to define addresses.
  - Blocked (IPS top attackers, top blocked/unblocked signatures, top signatures, top victims)—Whether the event resulted in dropped traffic (Blocked), traffic that was not dropped (Unblocked), or either (All).

- Username (Connection Profile Report and User Report)—The names of the users to include in the report. The default, an empty list, includes all users. If you want the report to focus on specific users, use the Add (+), Edit (pencil), and Delete (trash can) buttons beneath the table to create the desired list of users. Beginning with Version 4.7 of Security Manager, the Username filter supports uppercase/lowercase character sensitivity, wildcard characters, and "NOT" operation support. Please refer to **Enabling Uppercase/Lowercase Character Sensitivity** and **Disabling Wildcard Support** in the paragraphs at the end of this topic.
- Technology (Connection Profile Report, User Report, and VPN Device Usage Report)—The type of remote access technology to include in the report: All, Clientless (SSL VPN), Full Client (SSL VPN), IPsec RA (IPsec remote access VPN).
- Connection Profile (Connection Profile Report)—Allows you to customize a Connection Profile Report in Report Manager by adding or editing a connection profile for a Clientless (SSL VPN), Full Client (SSL VPN), or IPsec RA (IPsec remote access VPN) topology. For more information, see Connection Profiles Page, page 31-8 and Configuring Connection Profiles (ASA, PIX 7.0+), page 31-7. Beginning with Version 4.7 of Security Manager, the Connection Profile filter supports uppercase/lowercase character sensitivity, wildcard characters, and "NOT" operation support. Please refer to **Enabling Uppercase/Lowercase Character Sensitivity** and **Disabling Wildcard Support** in the paragraphs at the end of this topic.
- Group Policy (Connection Profile Report)—Allows you to customize a Connection Profile Report in Report Manager by specifying a single or multiple group polices in the filter to generate the report for the user who has logged with the specified group policies. For more information, see Configuring Group Policies for Remote Access VPNs, page 31-26. Beginning with Version 4.7 of Security Manager, the Group Policy filter supports uppercase/lowercase character sensitivity, wildcard characters, and "NOT" operation support. Please refer to **Enabling Uppercase/Lowercase Character Sensitivity** and **Disabling Wildcard Support** in the paragraphs at the end of this topic.
- User Session Duration (Connection Profile Report and User Report)—As the title indicates, allows you to specify the duration of a user session as <= or >= a given number of hours.

Step 5 Click **OK** in the Edit Settings dialog box to implement your changes.

You can now click the Generate Report button to retrieve the data defined by the settings and display it in the report. You can also use Save or Save As to permanently save the changes to the settings.

Enabling Uppercase/Lowercase Character Sensitivity

Beginning with Version 4.7 of Security Manager, uppercase/lowercase character sensitivity is available for the Username filter, the Connection Profile filter, and the Group Policy filter. Uppercase/lowercase character sensitivity is disabled by default. If you want to use it, you must enable it by following these steps; note that enabling it for one of these three filters enables it for all three of them:

1. Locate the reporting.properties file. The default location is NMSROOT/MDC/reports/config. (The default value for NMSROOT is C:\Program Files (x86)\CSCOpx
2. Set the parameter reports.reportgeneration.vpnUserReport.casesensitive.enable=true
3. Restart the Report Manager service, which is CsmReportServer

Disabling Wildcard Support

Beginning with Version 4.7 of Security Manager, wildcard support is available for the Username filter, the Connection Profile filter, and the Group Policy filter. Wildcard support is enabled by default. If you do not want to use it, you must disable it by following these steps; note that disabling it for one of these three filters disables it for all three of them:

1. Locate the reporting.properties file. The default location is NMSROOT/MDC/reports/config. (The default value for NMSROOT is C:\Program Files (x86)\CSCOpx

2. Set the parameter reports.reportgeneration.vpnUserReport.wildcard.enable=false
3. Restart the Report Manager service, which is CsmReportServer

Drilling Down into Report Data

The Top Destinations, Top Services, and Top Sources firewall reports and the Top Attackers, Top Signatures, and Top Victims IPS reports allow you to drill down into the report data.

To drill down into one of the supported reports, click on a data point in the Pie, XY, or Bar graph for that report. For example, if you click on a slice in the Pie chart for the Top Signatures report, you will see report details for the top attackers and top victims for the selected signature.

Note The filter criteria you use for a report affects the data available in the associated drill-down reports. If you apply filters to the report data, only one drill-down report will be shown when drilling down into the report data.

The following table shows the drill-down reports that are shown for each supported report type.

Report Name	Drill-down Reports Shown
Firewall	
Top Sources	Top Destinations, Top Services
Top Destinations	Top Sources, Top Services
Top Services	Top Sources, Top Destinations
IPS	
Top Signatures	Top Attackers, Top Victims
Top Attackers	Top Signatures, Top Victims
Top Victims	Top Signatures, Top Attackers

For each of the drill-down reports, you can view a chart or a table of the drill-down data. If you have drilled down into a specific data point in one of the supported reports, the drill-down report graph and tabular data will be included in the printed report and in the exported report data.

The following illustration shows an example of the drill-down report data for the Top Attackers report.

Figure 70-5 Top Attackers Drill-down Reports

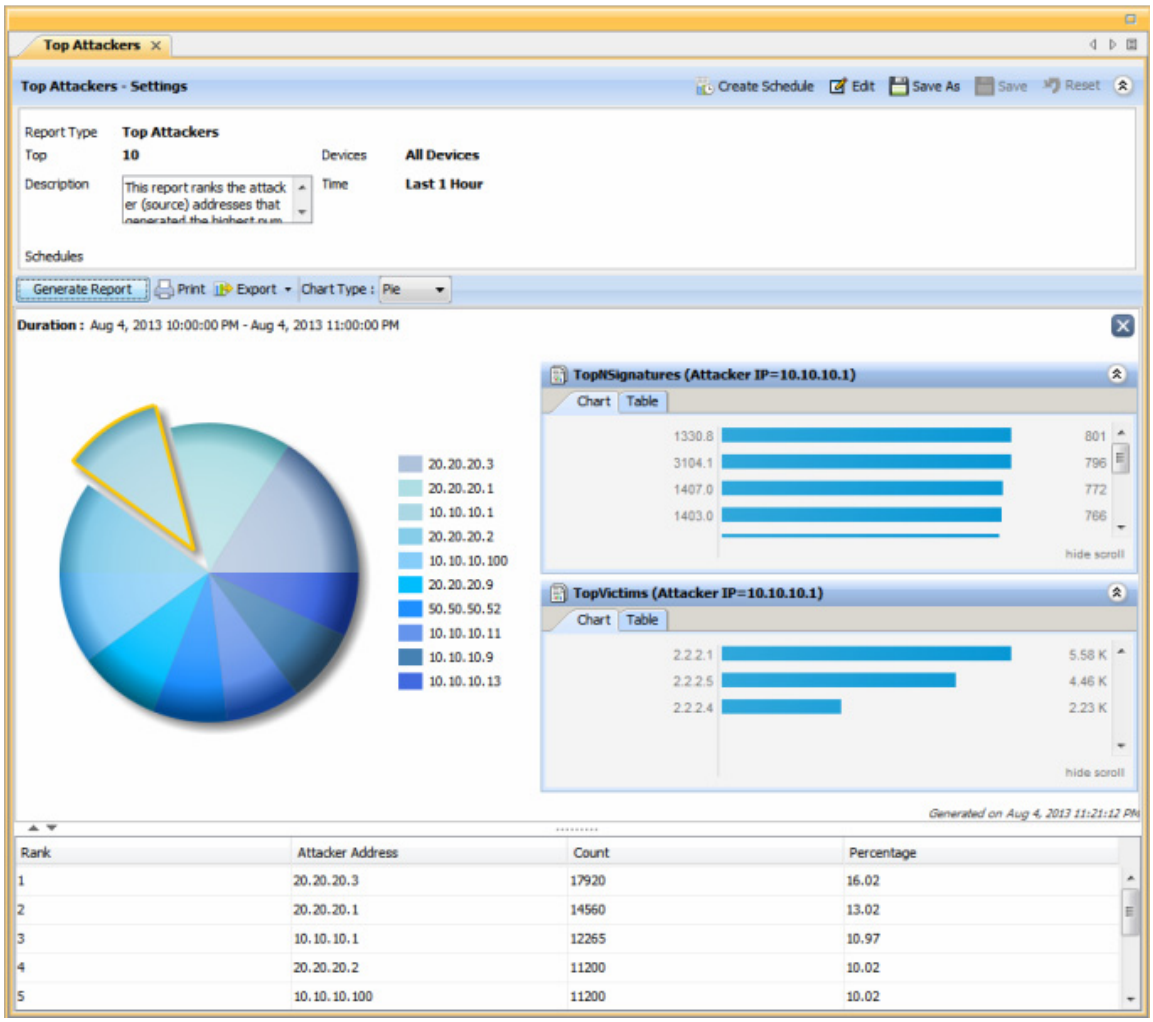

Related Topics

- Opening and Generating Reports, page 70-20
- Overview of Report Manager, page 70-6

Printing Reports

After you generate a report as described in Opening and Generating Reports, page 70-20, you can print it.

Note If you have opened drill-down reports for a specific data point in the Top Destinations, Top Services, or Top Sources firewall reports or the Top Attackers, Top Signatures, or Top Victims IPS reports, the drill-down report graph and tabular data will be included in the printed report. For more information, see Drilling Down into Report Data, page 70-26.

To print the report, click the **Print** button above the report. You are prompted to select a printer.

Exporting Reports

After you generate a report as described in Opening and Generating Reports, page 70-20, you can export it to an Adobe Acrobat (PDF) or comma-separated values (CSV) file.

Exported files include the following information:

- The creation time of the report.
- The settings used to generate the report.
- (PDF only.) The graphical representation of the report data.
- The tabular report data. In PDFs, the information is represented as a table. In CSVs, the information is comma-separated, with the first row being the column headings.

Note If you have opened drill-down reports for a specific data point in the Top Destinations, Top Services, or Top Sources firewall reports or the Top Attackers, Top Signatures, or Top Victims IPS reports, the drill-down report graph and tabular data will be included in the export. For more information, see Drilling Down into Report Data, page 70-26.

To export the report, click the down arrow in the **Export** button above the report and select either **As PDF** or **As CSV**. You are prompted to select a folder for the report. A default file name is provided, but you can change the file name.

Troubleshooting Report Exporting Error

When we try to export the device status report, we get the following error.

"Windows cannot find 'acrord32'. Make sure you typed the name correctly and then try again".

First, the above error occurs if Adobe Reader is not installed on the server. Windows cannot locate the acrord32.exe file as Adobe is not installed.

Second, even if Adobe reader is installed, the above error might be thrown.This is an issue which exists in windows XP,Vista, 7,8.1 and 10. This is caused due to failure in opening Adobe Reader. It is a known error and can happen to any application and not just Adobe reader. Microsoft has not provided patch for this yet.

Reported reasons why this issue can happen are:

1) Corrupt registry entry

2) Problem while installing Adobe

3) Deleting the default Adobe Reader

Following is the recommended workaround:

- If Adobe reader is not installed, install it
- If Adobe reader is installed and error is still thrown, Go to the location where the file is saved and perform Open with Adobe Reader action. Even though the error is thrown, the file will be created but it will be without a format. So we can open it using a PDF reader.

Configuring Default Settings for Reports

You can control the default settings that Report Manager uses for System reports. When you change the defaults, you automatically change the settings for all System reports, but the changes do not apply to any reports you saved as custom reports (in the My Reports folder). You must have system administrator or network administrator privileges to change these settings.

While viewing a report, you can edit any system report to specify different values for these settings, either to use temporarily while viewing the report, or to save as a custom report in the My Reports folder.

Note When creating reports in My Reports/Custom Reports, ensure the required filters are applied at the time of report creation because filters cannot be applied after the reports are created.

Tip If you change the settings in a system report, you can return the report to the default settings by clicking the **Reset** button in the Report Settings toolbar.

Step 1 In Report Manager, select **Tools > Default Report Settings** to open the Default Report Settings dialog box.

Step 2 Configure any of the following options:

- **Top**—The number of results listed in any of the "Top" reports, which list the most-often occurring items of the type targeted by the report. For example, if you select 20, the firewall Top Destinations report shows the 20 most often occurring traffic destinations in events reported to Security Manager. The default is 10.

  If you select more than 10, all items after the tenth are summarized as "others" in charts, although the detailed information for each additional item is listed in the report table.

- **Time Range**—The time window for events included in the report:
  - Last 1 Hour—The last full one hour on the zeros, for example, if the current time is 11:45 AM, the Last 1 Hour report shows data from 10:00 to 11:00.
  - Last 1 Day—The last full day, from midnight to midnight. For example, if the current day is Tuesday, the Last 1 Day report shows data from Monday.
  - Last 1 Week—The previous Monday through Sunday.
  - Last 1 Month—The previous month. For example, if the current date is September 29, the Last 1 Month report shows data from August.

  For more information on the implications of the time range setting, see Understanding Report Manager Data Aggregation, page 70-4.

- **Default Email Address**—The email address used as the default destination for scheduled reports.

If you decide you want to return to the installation default values, click **Restore Defaults**.

Step 3 Click **Apply** to save your changes, then click **Close** to close the dialog box.

Arranging Report Windows

You can open up to five report windows at one time. Reports are opened as tabbed windows in the right pane of the main Report Manager window, in the most recently used area ("tabbed group") if there is more than one area. The commands to arrange the windows appear if you right-click the tab for the report window as shown in the following illustration.

You have many options for arranging report windows based on your requirements. For example, you might want to compare two reports side-by-side, or remove a report from the main window without closing it.

You can use the following techniques to arrange the report windows to get the view that you desire:

- Floating a report—To remove a report from the main Report Manager window without closing it, right-click the report tab and select **Floating**. The report is moved to its own window.

  If you have already floated a report, you can select **Floating to** and choose one of the already-floated windows. The report becomes a new tab in that window.

- Docking a report—To move a floating report back to the main Report Manager window, right-click the report tab and select **Docking**.

- Arranging reports horizontally or vertically for side-to-side comparison—To create a vertical or horizontal arrangement of reports to allow for easy comparison, without floating the reports, right-click the report tab and select **New Horizontal Group** or **New Vertical Group**. These commands split the current tabbed group into the selected layout. You must have at least two reports open to use these commands. If you have more than two open reports, and you want all of them in separate windows, you need to use the command multiple times.

- Move reports to different tabbed groups—If you have several open reports, and you have arranged them into horizontal or vertical groups, you can move reports among the groups by right-clicking the report tab and selecting **Move to Next Tab Group** or **Move to Previous Tab Group**. The commands appear only if reports are arranged in a manner where such movement is possible:

- Change the orientation of groups—You can switch between horizontal and vertical layouts by right-clicking the report tab and selecting **Change Tab Groups Orientation**.

Saving Reports

If you edit the settings for a report, you must save it to make those changes permanent. However, to save changes to predefined system reports, you must save the report as a custom report.

Tip When you save a report, you are saving the settings that define the report. You are not saving the generated content of the report. If you want to save the generated content of the report, that is, the graphs and the report data, you must export the report rather than save it.

- To save changes to a custom report, do one of the following in Report Manager:
  - Select **File > Save** from the menu bar.
  - Click the **Save** button in the Report Settings toolbar.
- To save your changes as a new custom report, do one of the following to open the Save Report As dialog box:
  - Select **File > Save As** from the menu bar.
  - Click the **Save As** button in the Report Settings toolbar.
  - Right-click the report in the reports list and select **Save As**.

  Then, enter a name for the report and optionally a description of the report and click **OK**. The report is added to the My Reports folder in the reports list.

Note Report names can be up to 64 characters and contain alphanumeric characters, spaces, hyphens (-), and the underscore character (_). The description can be up to 1024 characters.

Renaming Reports

You can change the name of custom reports, but you cannot change the names of predefined system reports.

Related Topics

- Overview of Report Manager, page 70-6
- Understanding the Report List in Report Manager, page 70-9

Step 1 In Report Manager, select the report whose name you want to change from the reports list. You do not need to open the report; simply select it in the list.

Step 2 Click the **Edit** (pencil) button above the reports list to open the dialog box that allows you to change the report name.

Step 3 Enter the new name and click **OK**.

Report names can be up to 64 characters and contain alphanumeric characters, spaces, hyphens (-), and the underscore character (_).

Closing Report Windows

You can close any report by clicking the X icon on the report tab. For reports that you have floated, you can simply click the X icon in the window's title bar.

Tip When you close a report, none of the generated report data is preserved. If you want to preserve the generated data, you must print it or export it before you close the report window.

You can also use the following techniques to close the report windows without exiting Report Manager:

- Close a report—Select **File > Close Report** to close the currently-viewed report, or right-click the desired report tab and select **Close**.
- Close all reports—Select **File > Close All Reports**, or right-click any report tab and select **Close All**.
- Close all reports except one—Right-click the report tab of the report you want to keep open and select **Close Others**.

Deleting Reports

You can delete custom reports, but you cannot delete predefined system reports.

To delete a custom report, select it in the reports list and click the **Delete** (trash can) button above the reports list. You are asked to confirm your deletion.

Tip Deleting a report also deletes any schedules for that report.

If you need to delete another user's custom report, see Managing Custom Reports, page 70-32.

Managing Custom Reports

If you have system administrator or network administrator privileges, you can view a list of custom reports created by all Report Manager users on this Security Manager server.

To view the list of custom reports, select **Tools > Custom Report List** to open the Manage Custom Reports dialog box. The list shows the report name, the type of report, the type of device analyzed by the report, and the username of the person who created the custom report.

You can use the following controls to manage custom reports from this page:

- **Pagination controls**—If there are a lot of custom reports, use the pagination controls to move through the list. You can click the buttons to move to the first page, previous page, next page, or last page, or type the page number into the Page X of Y edit box. You can also click the down arrow in the edit box to change the edit box to work by record number rather than page number.
- **Delete button**—Click this button to delete the selected report. Any schedules (and schedule results) that use the report are also deleted.
- **Refresh button**—Click this button to update the list with the latest information.

Scheduling Reports

You can create schedules to regularly generate reports from Report Manager.

This section contains the following topics:

- Viewing Report Schedules, page 70-33
- Configuring Report Schedules, page 70-33
- Viewing Scheduled Report Results, page 70-35
- Enabling and Disabling Report Schedules, page 70-35
- Deleting Report Schedules, page 70-36

Viewing Report Schedules

You can view a list of report schedules that are configured in Report Manager. If you have system administrator or network administrator privileges, the list includes all schedules configured on the server, whether you configured them or another user configured them. Users with lesser privileges can see their own schedules only.

To view the list of report schedules, select the **Scheduled Reports** tab, and if necessary, the **Schedule List** sub-tab. The list shows the schedule name, description, the report that will be generated by the schedule, the frequency of report generation, the e-mail addresses to which reports are sent (if any), whether the schedule is enabled or disabled, and the username of the person who created the schedule.

You can use the following controls to manage report schedules from this page:

- **Pagination controls**—If there are a lot of schedules, use the pagination controls (below the table to the left) to move through the list. You can click the buttons to move to the first page, previous page, next page, or last page, or type the page number into the Page X of Y edit box. You can also click the down arrow in the edit box to change the edit box to work by record number rather than page number.
- **Add button**—Click this button to add a new schedule. For more information, see Configuring Report Schedules, page 70-33.
- **Edit button**—Click this button to edit the selected schedule. For more information, see Configuring Report Schedules, page 70-33.
- **Delete button**—Click this button to delete the selected schedule. For more information, see Deleting Report Schedules, page 70-36.
- **Refresh button**—Click this button to update the list with the latest information.
- **Enable button**—Click this button to enable the selected schedule. The button is active only if the selected schedule is disabled. For more information, see Enabling and Disabling Report Schedules, page 70-35.
- **Disable button**—Click this button to disable the selected schedule. The button is active only if the selected schedule is disabled. For more information, see Enabling and Disabling Report Schedules, page 70-35.

Configuring Report Schedules

You can create schedules to automatically generate reports at set times. The generated reports are e-mailed to identified recipients and are also stored and available for viewing from Report Manager. By scheduling reports, you can easily and efficiently create regular milestone views of network security and usage. This procedure explains how to set up a schedule for a report.

Related Topics

- Overview of Report Manager, page 70-6
- Opening and Generating Reports, page 70-20
- Viewing Report Schedules, page 70-33
- Troubleshooting Report Manager, page 70-36

Step 1 In Report Manager, do one of the following:

- On the Reports tab, open the report for which you are creating a new schedule by double-clicking the name of the report in the report list (in the left pane). Then, click the **Create Schedule** button in the report settings toolbar.

  Note You cannot edit an existing schedule from the Reports tab.

- On the Reports tab, right-click the report on which you are creating a schedule and select **Create Schedule**. If the report is not already open, it is opened.
- On the Scheduled Reports tab, Schedule List sub-tab, click the **Add** button below the list of schedules to create a new schedule. To edit an existing schedule, select it in the list and click the **Edit** button.

The Add or Edit Report Schedule dialog box opens.

Step 2 Configure the following options in the dialog box:

- **Schedule Name**—A name for the schedule, up to 64 characters.
- **Report Name**—Select the name of the report to be generated by the schedule. When you create a schedule from the report settings pane, the name is pre-selected and you cannot change it.
- **Schedule**—Select whether the report shall be generated daily, weekly (once per week), or monthly (once per month). Then, enter the day and time for generating the report:
  - Daily schedules—Select whether the schedule is for Monday through Friday (five days) or Monday through Sunday (all seven days). Enter the time of day (in 24-hour notation) to generate the report.
  - Weekly schedules—Select the day of the week and enter the time of day in 24-hour notation.
  - Monthly schedules—Select whether to generate the report on the first day of the month, the last day, or custom. If you select Custom, enter the day number. Then, enter the time of day in 24-hour notation.
- **Email To**—The email addresses that should be sent the report. Separate multiple addresses with commas. If you do not want reports e-mailed, ensure that the field is empty. Note that for e-mails to be sent successfully, you must configure SMTP on the Security Manager server as described in Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.

  If the report fails to generate for some reason, notification of the failure is sent to these e-mail addresses.
- **Export Report Format**—Whether to generate the report in Adobe Acrobat (PDF) or comma-separated value (CSV) format. The PDF includes graphics, the CSV does not. For more information on export formats, see Exporting Reports, page 70-28.
- **Description**—A description of the schedule.
- **Status**—Whether the schedule is enabled (reports will be generated) or disabled (reports will not be generated).

Step 3 Click **OK** to save the schedule. New schedules are added to the schedules list on the Schedules tab.

Viewing Scheduled Report Results

Typically, report schedules include e-mail addresses to which generated reports are sent. You can also view reports generated from schedules in Report Manager. If you have system administrator or network administrator privileges, you can view results generated by other users' schedules.

Tip Report Manager maintains a copy of the last report generated by the schedule. You cannot retrieve previously generated reports.

Related Topics

- Overview of Report Manager, page 70-6
- Opening and Generating Reports, page 70-20
- Viewing Report Schedules, page 70-33
- Troubleshooting Report Manager, page 70-36

Step 1 In Report Manager, select the **Scheduled Reports** tab.

Step 2 Select the **Results** sub-tab.

All results that you are authorized to see are listed on this tab. The list shows the schedule name, the name of the report that was generated, the frequency of report generation, the date and time of the last schedule run (when the report was generated), the status of the report generation (Success or Failed), a link to the generated report (in the Last Report column), and the username of the person who created the schedule.

Tip If the status of a report is Failed, click the link to see the reason for failure.

Step 3 Double-click the icon link to the report in the Last Report column to open it. While viewing the report, you can save it to your workstation.

If you cannot find the report you are looking for, click the **Refresh** button to update the list with the latest information.

Enabling and Disabling Report Schedules

You can enable or disable report schedules to change whether reports are generated based on the schedules. By disabling a schedule, you can prevent reports from being generated without deleting the schedule. If you have system administrator or network administrator privileges, you can enable or disable another user's schedule.

Related Topics

- Overview of Report Manager, page 70-6

- Viewing Report Schedules, page 70-33

Step 1 In Report Manager, select the **Scheduled Reports** tab, then if necessary, the **Schedule List** sub-tab. This tab lists all currently-defined schedules that you are authorized to see.

Step 2 Select the schedule whose status you want to change and click either the **Enable** or **Disable** button.

Deleting Report Schedules

You can delete report schedules if you no longer need them. If you have system administrator or network administrator privileges, you can delete another user's schedule.

Tip If you do not want to generate reports from a schedule, but you want to keep the schedule definition, you can disable the schedule. Disabled schedules do not generate reports.

Step 1 In Report Manager, select the Scheduled Reports tab, then if necessary, the Schedule List sub-tab. This tab lists all currently-defined schedules that you are authorized to see.

Step 2 Select the schedule and click the **Delete** button below the list. You are asked to confirm the deletion.

When you delete a schedule, any results of the schedule are also deleted from the server and they are removed from the Results tab.

Troubleshooting Report Manager

Following are some problems you might encounter when using the Report Manager application, and some ways to resolve the problems.

Problem: Report Manager does not open, you get the message "Not able to connect to server."

Solution: Report Manager requires that the csmReportServer, rptDbEngine, and rptDbMonitor processes be started. Report Manager also relies on the Event Management service VmsEventServer. Ensure that all services are started and running correctly on the Security Manager server.

To view the current state of the processes, log into the Security Manager web interface using http://*SecManServer*:1741, where *SecManServer* is the DNS name of the server. From the Security Management Suite home page, click the **Server Administration** link to open CiscoWorks Common Services at the Admin page. Click **Processes** in the TOC on the left side of the window to open the list of processes that displays their current states. Select these processes and click **Start** to start them. If necessary, you might want to stop them, and then restart them. Wait for the processes to fully restart, then try again to open Report Manager.

Problem: When generating a report, you get the message "No records found."

Solution: This message indicates that there were no event records in the event data storage location that relate to the report type and the configured settings, or that the required Report Manager aggregation cycle has not completed. Investigate the following:

- Ensure that devices of the appropriate type are selected for monitoring as described in Selecting Devices to Monitor, page 69-34.

- Ensure that these devices are appropriately configured for sending events to Security Manager, and that events from the device appear in Event Viewer. Ensure that the device and Security Manager are using the same syslog port. For information on configuring the devices, see Configuring ASA and FWSM Devices for Event Management, page 69-28 and Configuring IPS Devices for Event Management, page 69-29. To check the syslog port that Security Manager is using, view the setting on the **Tools > Security Manager Administration > Event Management** page in Configuration Manager.
- For IPS devices, ensure that the certificate has not expired. Check the certificates table by selecting **Manage > IPS > IPS Certificates** in Configuration Manager and regenerate the certificate if necessary.
- The report settings might specify a time period in which no aggregated data exists for the report. Data is aggregated every 15 minutes, 1 hour (on the hour), and 1 day (at midnight). Try changing the time parameters of the report. See Editing Report Settings, page 70-22. Consider the following:
  - To see a Last 1 Hour report, a change of hour must occur since initially starting the Event Manager service. For example, if you start the service at 10:05, hourly reports are available only after 11:00.
  - To see a Last 1 Day report, a change of day must occur since initially starting the Event Manager service. For example, if you start the service at 10:05, you must wait until after midnight to see a daily report.
  - To see a Last 1 Week report, at least one full day cycle must occur. Weekly reports are based on daily reports.
  - To see a Last 1 Month report, at least one full month must pass since starting the service.
  - To see a custom time period report, at least one daily cycle must occur.
- If you create a new custom report, it might take up to one hour for data to become available. Also, ensure that the time period is Last 1 Hour until the report is old enough to have data for other time periods.

Problem: You cannot get reports for a specific device.

Solution: Investigate the following:

- The device must have been selected for event management during the reporting time-frame as described in Selecting Devices to Monitor, page 69-34. Even if a device is selected, Report Manager might not support all devices that are supported for Event Viewer. For information on the supported device types, see Understanding Report Management, page 70-1.
- The report settings might exclude the device. Unless the report settings indicate that All Devices are considered in the report, check the device selection to ensure the device is included. See Editing Report Settings, page 70-22.
- The report settings might specify a time period in which no data exists for the device. Try changing the time parameters of the report. See Editing Report Settings, page 70-22.
- If your organization uses Cisco Secure ACS to control access to the application, you can view reports on a device only if you have at least View privileges to the device. Check whether you have the required permissions.

Problem: You cannot see data in certain IPS predefined reports after specifying values for each of signature, victim IP, and attacker IP.

Solution: The top attackers, top victims, and top signatures predefined reports include criteria for signature, victim IP address, and attacker IP address. However, you cannot configure all three criteria in a predefined report. Instead, you can configure the criteria on which the report is based (for example, victim IP address for the top victims report) plus one, and only one, of the remaining values. Note that this limitation does not apply to the other criteria, such as Blocked and Top.

Problem: You cannot see data in certain Firewall predefined reports after specifying values for each of service, source IP, and destination IP.

Solution: The top destinations, top services, and top sources predefined reports include criteria for service, source IP address, and destination IP address. However, you cannot configure all three criteria in a predefined report. Instead, you can configure the criteria on which the report is based (for example, service for the top services report) plus one, and only one, of the remaining values. Note that this limitation does not apply to the other criteria, such as Permit/Deny and Top.

Problem: Statistics for VPN reports are not available.

Solution: VPN statistics are partially obtained directly from the device rather than from events stored in the event data storage location. To obtain the statistics, Report Manager must be able to log into the device and use show commands. Ensure that the device properties for your VPN devices have the correct credentials for logging in.

Problem: Scheduled reports are not getting sent to recipients.

Solution: Ensure that the SMTP server is configured correctly and that a valid source e-mail address is configured for Security Manager. For more information, see Configuring an SMTP Server and Default Addresses for E-Mail Notifications, page 1-27.

Problem When exporting the device status report, you get the following error.

```
"Windows cannot find 'acrord32'. Make sure you typed the name correctly and then try
again".
```

Solution Execute the following:

- Install Adobe Reader on the server, if it is not installed already. MS-Windows will not be able to locate the acrord32.exe file if Adobe Reader is not installed.
- Though Adobe Reader is installed, the error might be thrown if you are using Windows XP, Vista, 7,8.1 or 10. This is a known error of Microsoft Windows. Microsoft is yet to provide patch for this error. Do the following:
  - Go to the location where the exported report file is saved. Right-click and choose **Open with > Adobe Reader**. Despite the error, the file is created without the defined format. Hence, you can open it using a PDF reader.

Health and Performance Monitoring

The Health and Performance Monitor (HPM) application lets you monitor key health and performance data for ASA devices, IPS devices, and VPN services by providing network-level visibility into device status and traffic information.

A variety of views are provided—All Devices, Firewall Devices, IPS Devices, VPN Summary, and so on—and you can create your own customized views. A configurable listing of device alerts is also available.

This ability to monitor key network and device metrics lets you quickly detect and resolve device malfunctions and bottlenecks in the network.

This chapter contains the following topics:

- Health and Performance Monitor Overview, page 71-1
- HPM Access Control, page 71-3
- Preparing for Health and Performance Monitoring, page 71-4
- Launching the Health and Performance Monitor, page 71-4
- Managing Monitored Devices, page 71-5
- HPM Window, page 71-6
- Monitoring Devices, page 71-21
- Alerts and Notifications, page 71-32
- SNMP Trap Forwarding Notification, page 71-44

Health and Performance Monitor Overview

The Health and Performance Monitor is a stand-alone application that you can launch from the other stand-alone Security Manager applications (Dashboard, Configuration Manager, Event Viewer, Report Manager, and Image Manager) or from the Cisco Security Manager Client login screen accessed from the Windows Start menu.

The HPM application complements the Event Viewer and Report Manager applications, as follows:

- **Event Viewer** – Monitors your network for syslog (system log) events from ASA and FWSM devices and their security contexts, and for SDEE (Secure Device Event Exchange) events from IPS devices and virtual sensors. These events include firewall traffic information, NAT events, failover events, IPS alerts, and so on. Event Viewer collects and displays this information, organized into a variety of views. See Chapter 69, "Viewing Events" for more information.

- **Report Manager** – Collects, displays and exports network usage and security information for ASA and IPS devices, and for remote-access IPsec and SSL VPNs. These reports aggregate security data such as top sources, destinations, attackers, victims, as well as security information such as top bandwidth, duration, and throughput users. Data is also aggregated for hourly, daily, and monthly periods. See Chapter 70, "Managing Reports" for more information.
- **Health and Performance Monitor** (HPM) – Monitors and displays key health, performance and VPN data for ASA and IPS devices in your network. This information includes critical and non-critical issues, such as memory usage, interface status, dropped packets, tunnel status, and so on. You also can categorize devices for normal or priority monitoring, and set different alert rules for the priority devices.

  You can add notes to displayed alerts, you can "acknowledge" them, and you can clear them. When an alert is cleared, it is removed from the Alerts display; however, the alert information is retained in a database for 30 days. See Alerts: Acknowledging and Clearing, page 71-42 for more information about adding notes, and acknowledging and clearing alerts.

Note You can use the Alerts History window to access and view previously cleared alerts, as described in Alerts: History, page 71-43.

This section contains the following topics:

- Trend Information, page 71-2
- Monitoring Multiple Contexts, page 71-3

Trend Information

The Health and Performance Monitor periodically polls monitored devices for status and performance data. This information is used for alert generation, and to display real-time views and historical trends based on aggregated data.

Trends are displayed graphically for a specific set of metrics. Each trend for the currently selected device is represented as a graph generated for a chosen time interval. Comparing current values with the weekly averages for CPU and memory usage, for example, can provide an operational context for the selected device. Available trend intervals for monitored devices are one hour, 24 hours, and one week.

Metrics used for generating trends include:

- CPU usage
- Memory usage (only for single-context devices)
- Connections per second (firewall devices)
- Translations per second (firewall devices)
- Inspection load (IPS devices)
- Missed packets as a percentage (IPS devices)
- Number of VPN tunnels
- Number of RA VPN sessions
- Total VPN throughput
- Firewall throughput
- Total dropped packets (firewall interfaces)

For additional graphical information about the health and performance of a specific device, you can launch the related device manager by right-clicking the entry for a device, a cluster node, or the system context for a multi-context device, and then choosing **Device Manager** from the pop-up menu. See Starting Device Managers, page 72-14 for more information about the device managers.

Monitoring Multiple Contexts

The Health and Performance Monitor can monitor single- and multiple-context ASA devices. For multiple-context devices, each context is monitored and displayed as if it was a separate device.

Each context will be polled separately for all applicable metrics, with HPM polling a maximum of five contexts at a time from any given device. For devices with more than five contexts, data will be acquired from each successive batch of five contexts, with each batch being polled progressively during successive polling cycles. This means that all contexts may not be updated at the same time.

For multiple-context devices, basic device health—memory usage, device status, and so on—is monitored only on the physical device (that is, from the system context), while traffic data—number of connections, number of translations, dropped packets and so on—are monitored at context level.

For virtual contexts, CPU usage data are used only for pattern analysis, not for alert generation. Only interface-status alerts will be generated for virtual contexts.

HPM Access Control

The privileges assigned to your user name control what you can do in Health and Performance Monitor. If you use local users, or other types of non-ACS access control, then all users have access to HPM. However, the following access limits are imposed:

- You must have system administrator privileges to enable or disable Health and Performance Monitoring in Security Manager, as described in Health and Performance Monitor Page, page 11-36.
- You must have system administrator, network administrator, or approver privileges to select or deselect devices for monitoring, as described in Managing Monitored Devices, page 71-5.
- You also must have system administrator, network administrator, or approver privileges to configure alerts and notifications, as described in Alerts: Configuring, page 71-34.

If you use ACS to control access to Security Manager, you can also control the following:

- You can control access to the Health and Performance Monitor application using the View > Health and Performance Monitor privilege (part of Role Management in ACS). Using this privilege, you could prevent certain users from accessing HPM, or create roles that allow access to HPM without allowing access to Event Viewer or Report Manager. All default ACS roles are permitted to use the Health and Performance Monitor application.
- Use the Modify > Policies > HPM Monitoring privilege to control which users can select and deselect the devices that are monitored (see Managing Monitored Devices, page 71-5), configure alerts and notifications (see Alerts: Configuring, page 71-34), and annotate and acknowledge alerts (see Alerts: Acknowledging and Clearing, page 71-42). All default ACS roles except Help Desk and Super Admin have this permission.
- Users can view health and performance information for a device only if they have at least View privileges for the device.

- You can control access to the Health and Performance Monitoring administrative settings page (in Security Manager's Configuration Manager) where HPM is enabled or disabled, as described in Health and Performance Monitor Page, page 11-36. The user must have the Modify > Policies > HPM Admin privilege to access this page (or any other administrative settings page). All default ACS roles except Help Desk can view the page, but only System Administrators can change the setting.

For information on integrating Security Manager with Cisco Secure ACS, see the *Installation Guide for Cisco Security Manager*.

Preparing for Health and Performance Monitoring

In order to use the Health and Performance Monitor (HPM), you must configure Security Manager, enable the HPM application, and configure device monitoring, as follows:

- Basic Threat Detection must be enabled on ASA 8.0+ devices in order to monitor metrics such as ACL Dropped Packets, Scanning Threat Dropped Packets, Inspection Dropped Packets, and Syn Attack Dropped Packets. (Basic Threat Detection is enabled by default on these ASA devices.)
- To receive alert notifications via email, you must have configured the SMTP server and administrator email ID on the **System Preferences** page of the Security Manager server. See the *Installation Guide for Cisco Security Manager* for more information. (Specifying email addresses for alert notifications from the Health and Performance Monitor application is described in Alerts: Configuring, page 71-34.)
- Health and Performance Monitoring must be enabled in Security Manager, as described in Health and Performance Monitor Page, page 11-36.
- In HPM, specify the devices to be monitored, in both Normal and Priority modes, as described in Managing Monitored Devices, page 71-5.

Note To prevent read time-outs for ASAs, those devices must be configured to use only certain SSL/TLS protocol versions when acting as a server, as described in Setting Up SSL (HTTPS) on PIX Firewall, ASA and FWSM Devices, page 2-3.

- Enable and configure the device threshold values and state-change rules that define when alerts and email notifications are triggered. This process is described in Alerts: Configuring, page 71-34.

Note We also recommend configuring monitored devices to use a Network Time Protocol (NTP) server for synchronized timing. See NTP Page, page 52-21 for more information.

After you have completed these steps, HPM begins polling the specified devices and displays health information and alerts.

Launching the Health and Performance Monitor

Use the Health and Performance Monitor (HPM) to view status information and alerts collected from monitored firewall and IPS devices across your network. For more information about selecting devices for monitoring, see Managing Monitored Devices, page 71-5.

To launch HPM, do any one of the following:

- Choose **All Programs > Cisco Security Manager Client > Cisco Security Manager Client** from the Windows Start menu (your command path may differ slightly), and then select **Health and Performance Monitor** as the Default View during login.
- Choose **Launch > Health and Performance Monitor** from the Configuration Manager, Event Viewer, Image Manager, or Report Manager applications.
- Click the Health and Performance Monitor button on the quick-launch toolbar in the Configuration Manager or Image Manager window.

If you are currently not logged into a Security Manager application, you are prompted to log in. (For more information about starting and logging into a Security Manager client application, see Logging In to and Exiting the Security Manager Client, page 1-12). Otherwise, the HPM Window, page 71-6 is opened using the same user account you used to log into the other application.

Note As described above, you can "cross-launch" HPM from any of the other Security Manager client applications. You can similarly cross-launch any of the other client applications from Health and Performance Monitor by choosing the desired application from the **Launch** menu, or clicking the appropriate quick-launch button.

Managing Monitored Devices

The HPM device selector is used to add and remove devices from both the "normal" and "priority" monitoring lists. You can also use the device selector to transfer devices between the two lists.

Note After enabling a device for monitoring in HPM, it can take up to 5 minutes for priority devices and 10 minutes for non-priority devices before actual values for HPM parameters can be seen in the device summary.

To use the HPM device selector:

Step 1 Choose **Device Selector** from the Tools menu to open the device selector window; the device-management screen is displayed.

The All Devices section on the left lists all ASA and IPS devices in the Security Manager inventory that can be monitored. (For example. HPM supports monitoring of version 7.0.1 and later IPS sensors only. Earlier IPS versions are not displayed in the device selector.)

All devices currently assigned to the Normal monitoring list and the Priority monitoring list are displayed in the two sections on the right side of the window.

Step 2 To add a device to the Normal list, select the device in the All Devices list and then click the > button between the All Devices list and the Normal Monitored Devices list.

The procedure for moving a device to the Priority Monitored Devices list is the same: use the > button between the All Devices list and that list.

Step 3 To remove a device from either Monitored list, returning it to the All Devices list, select the device and then click the appropriate < button.

Step 4 To transfer a device from one Monitored list to the other, highlight that entry and click the Up or Down button to move it to the upper or lower list respectively.

Step 5 Click Next at the bottom of the window to display the VPN-selector screen.

All monitored devices and their individual contexts, if any, are listed; each entry includes a checkbox for remote-access (RA) and one for site-to-site (S2S) VPN selection.

Note Starting from Cisco Security Manager 4.10, all the contexts for ASA 9.5(2) and above will be listed in the **Device Selector**. From the **Device Selector**, you can now monitor the RA and site-to-site VPN for all user contexts by enabling the corresponding check box.

You can use the List Filter field on this page to filter the list, as described in Using The List Filter Fields, page 71-19.

Step 6 Select the types of VPN to be monitored on specific devices by checking the appropriate boxes.

Step 7 Click **Save** to save and apply your changes, and close the device selector.

HPM Window

The Health and Performance Monitor (HPM) application window is where you view status information and alerts collected from monitored firewall and IPS devices, as well as remote-access (RA) and site-to-site (S2S) VPN information, across your network.

Note See Managing Monitored Devices, page 71-5 for information about specifying the devices to be monitored.

The following illustration presents the primary features of the HPM window.

Figure 71-1 Health and Performance Monitor Window

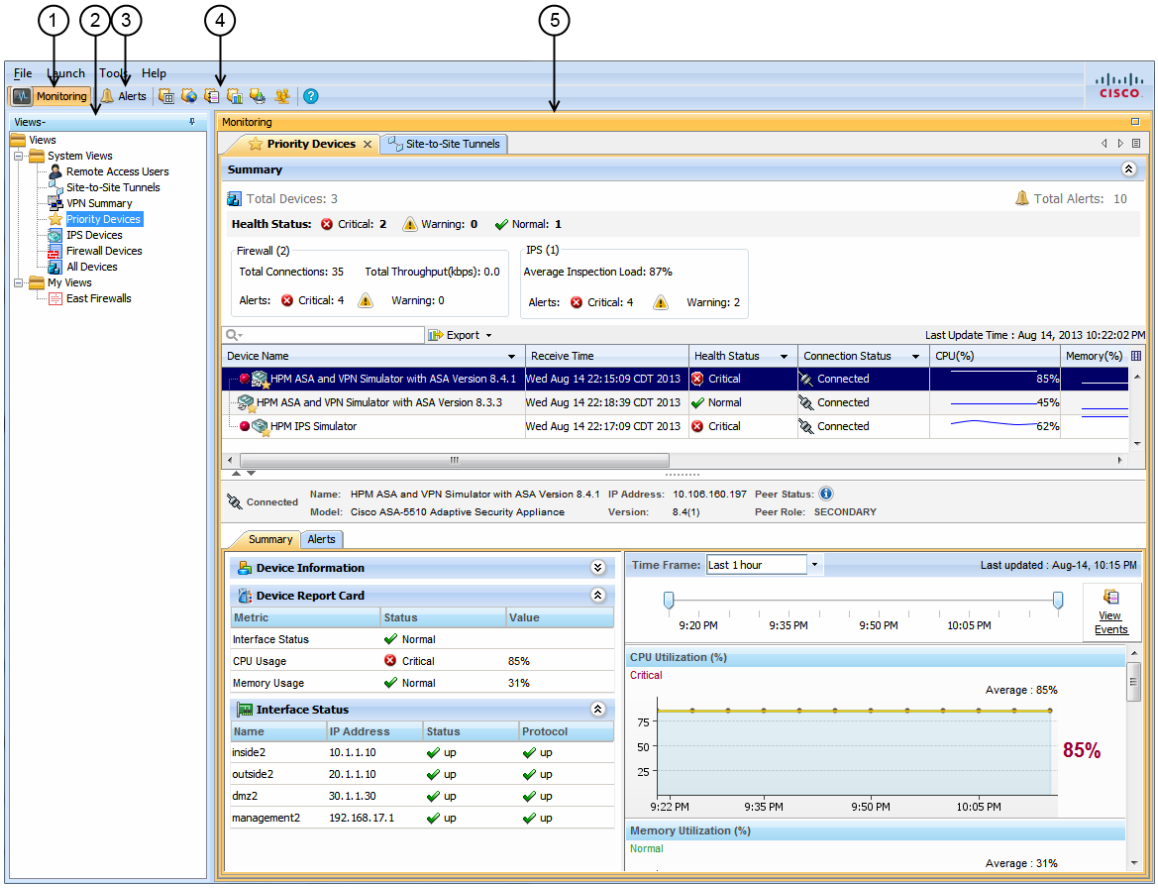

1	Monitoring button.	**4**	Quick-launch buttons.
2	Views	**5**	Monitoring/Alerts display area.
3	Alerts button.		

The HPM window consists of three main elements:

- **Monitoring button (1)** – Click this button to view device and VPN health and performance data. See HPM Window: Monitoring Display, page 71-25 for more information.
- **Views (2)** - When in the Monitoring view, the left pane of the HPM main window displays a list of available views. See Managing Device Views, page 71-21 for more information.

  as shown in the following illustration.
- **Alerts button (3)** – Click this button to view a table of alerts in the window's display area. See HPM Window: Alerts Display, page 71-32 for more information.
- **Quick-launch buttons (4)** – Click any button to cross-launch the related Security Manager client application.
- **Monitoring/Alerts display area (5)** – This section of the window displays either Monitoring information for devices and VPNs, or a table of alerts generated by monitored devices. The Monitoring and Alerts buttons are used to switch back and forth between these two displays.

Working with Table Columns

You can customize the different tables of information presented in HPM as follows:

- Sort a table such entries in a particular column are in ascending or descending order.
  - Click a column heading—anywhere but on a drop-down menu button—to sort the table such that the column entries are in ascending order (indicated by a small grey up-arrow).
  - Click the heading again to sort the entries are in descending order (indicated by a small grey down-arrow).
  - Click the heading again to return the table to its original order of display (the direction icon is removed).
- Hide and show various columns; the columns available for display depend on the particular table.
- Apply a column filter, meaning the table displays only entries that fit the specified criteria.

This section contains the following topics:

- Showing and Hiding Table Columns, page 71-8
- Column-based Filtering, page 71-17

Showing and Hiding Table Columns

You can customize the different tables presented in HPM by hiding and showing various columns of information; the columns available for display depend on the particular table.

Note The column headings are menus that you can use to further filter the table by hiding or showing entries according to chosen parameters, as described in Column-based Filtering, page 71-17.

To show or hide specific columns displayed for a table:

1. Click the Columns button on the right side of the column headings to open the Choose Columns to Display dialog box.

   All columns available for the current view are listed.
2. Select and deselect the columns to be shown and hidden.
3. Click **OK** to close the dialog box.

   Only the selected columns are displayed for this table.

The following topics describe the individual columns available for various tables:

- Table Columns: Device-related Views, page 71-8
- Table Columns: VPN-related Views, page 71-13
- Alert Table Columns, page 71-16

Table Columns: Device-related Views

You can customize the tables presented in the Monitoring pane for the device-related views by hiding and showing various columns of information; the columns available for display depend on the particular view.

The order of the entries in the Choose Columns to Display dialog box reflects the ordering of the columns when displayed. (However, the ordering of the rows in the following table does not necessarily reflect ordering of the columns as displayed.) See Showing and Hiding Table Columns, page 71-8 for information about opening the Choose Columns to Display dialog box.

The following table presents all available data columns for the device-related Monitoring views: Priority Devices, IPS Devices, Firewall Devices, All Devices, and all custom views based on these system views. Some of the listed columns are not available for specific views, as indicated.

Table 71-1 Available Table Columns for Device-related Views

Column Name	Available in View*	Description
Device Name	IPS, Firewall	Name assigned to the device; that is, the Host Name as defined on the Device Properties: General Page, page 3-41 of the Device Properties window. Column-based Filtering, page 71-17 is available. The Name is preceded by an icon indicating device type. This icon in turn may be preceded by a device-alerts indicator: a red dot indicates one or more critical alerts (and possibly warnings), while a yellow dot indicates one or more warnings only. The area is blank for a device with no alerts. You can "hover" the mouse pointer over the dot to view a pop-up displaying the number of critical alerts and the number of warnings on the device. A gold star is added to the device icon itself to indicate Priority monitoring.
Receive Time	IPS, Firewall	Poll date and time for this entry (format is: day-of-week MMM DD HH:MM:SS your-time-zone YYYY).
IP Address	IPS, Firewall	IP address of this device. Column-based Filtering, page 71-17 is available.
Health Status	IPS, Firewall	Current overall health of the device: Critical, Warning, or Normal. Column-based Filtering, page 71-17 is available. **Note** Overall health is defined by the most critical of any of the health metrics. For instance, if all the selected metrics on the device are normal except for one that is critical, overall device health becomes critical.

Table 71-1 Available Table Columns for Device-related Views (continued)

Column Name	Available in View*	Description
Connection Status	IPS, Firewall	Indicates HPM's ability to connect to/poll the device: Connected, Authentication Error, Certificate Mismatch Error, Connection error, Timeout during Read operation, or Service unavailable. Column-based Filtering, page 71-17 is available. **Note** If the device is not selected as a Normal or Priority Monitored Device in HPM (Tools > Device Selector), this status will not apply. Changes to Monitored Device selection may take several minutes to become effective and be reflected on screen. Any information displayed for a non-"Connected" device is from the indicated Receive Time, prior to connection failure.
Memory (%)	IPS, Firewall	Memory usage as a percentage of the total available.
CPU (%)	IPS, Firewall	CPU usage as a percentage of the total available.
Model	IPS, Firewall	Device type and model number. For example, ASA 5510, or IPS 4270.
Version	IPS, Firewall	Software version running on this device. Column-based Filtering, page 71-17 is available.
Inspection Load (%)	IPS	Inspection load on the device when polled, as a percentage.
Missed Packet(%)	IPS	Dropped packets as a percentage of total packets inspected.
Sensor App Status	IPS	Current Sensor App (Analysis Engine) status: Up or Down. Column-based Filtering, page 71-17 is available.
Main App Status	IPS	Current Main App status: Up or Down. Column-based Filtering, page 71-17 is available.
Collaboration App Status	IPS	Current Collaboration App status: Up or Down.
License Expiration Status	IPS	Status of the sensor's license, based on red and yellow threshold values set on the sensor: Normal, Warning, or Critical. Column-based Filtering, page 71-17 is available.
In Bypass Mode	IPS	Whether bypass mode is enabled on the sensor: Yes or No. Column-based Filtering, page 71-17 is available.
Event Retrieval Status	IPS	Status of the IPS event retrieval: Normal, Warning, or Critical. Column-based Filtering, page 71-17 is available.

Table 71-1 Available Table Columns for Device-related Views (continued)

Column Name	Available in View*	Description
Global Correlation Status	IPS	For a sensor participating in global correlation, its update status: Normal (last update was successful), Warning (no successful update within the past day [86,400 seconds]), or Critical (no successful update within the last three days [259,200 seconds]). Column-based Filtering, page 71-17 is available.
Signature Update	IPS	The number of the most recent signature update applied to this sensor; for example, S574. Column-based Filtering, page 71-17 is available.
Firewall Mode	Firewall	Operating mode of this device: Routed, Transparent, or Mixed. Column-based Filtering, page 71-17 is available.
Context Mode	Firewall	Context mode of this device: Single or Multiple. Column-based Filtering, page 71-17 is available.
Connections	Firewall	Number of active connections when device was polled.
Xlates	Firewall	Address translation counter.
Connections/second	Firewall	Number of connections established per second.
Translations/second	Firewall	Number of translations per second.
Failover Status	Firewall	If this device is part of a failover pair, its current state: Active or Standby. Column-based Filtering, page 71-17 is available.
Failover Host Role	Firewall	If this device is part of a failover pair, its current role: Primary or Secondary. Column-based Filtering, page 71-17 is available.
Failover Peer Role	Firewall	If this device is part of a failover pair, current role of its peer device: Primary or Secondary. Column-based Filtering, page 71-17 is available.
Failover Peer Status	Firewall	If this device is part of a failover pair, current status of its peer: Active or Standby Ready. Column-based Filtering, page 71-17 is available.
Used Memory (MB)	Firewall	Amount of memory (in megabytes) in use when device was polled. Column-based Filtering, page 71-17 is available.
Free Memory (MB)	Firewall	Amount of memory available (in megabytes) when device was polled. Column-based Filtering, page 71-17 is available.
Max. Connections	Firewall	Peak number of connections. Not available for ASA clusters.
Max. Xlates	Firewall	Peak number of address translations. Not available for ASA clusters.

Table 71-1 Available Table Columns for Device-related Views (continued)

Column Name	Available in View*	Description
Throughput (Kbps)	Firewall	Average device throughput in kilobits per second. For an ASA 9.0+ cluster, this the total throughput for all interfaces in the cluster.
ACL Dropped Packets	Firewall	The number of packets dropped because they failed an access control list rule. Available only at cluster level for ASA clusters; not available for individual nodes.
Scanning Threat Dropped Packets	Firewall	If scanning threat detection is enabled, the number of packets dropped because they failed scanning threat inspection. If not enabled, "NA" is displayed. Available only at cluster level for ASA clusters; not available for individual nodes.
Inspection Dropped Packets	Firewall	If application inspection is enabled, the number of packets dropped because they failed application inspection. If not enabled, "NA" is displayed. Available only at cluster level for ASA clusters; not available for individual nodes.
Syn Attack Dropped Packets	Firewall	Number of packets dropped because of SYN flooding. Available only at cluster level for ASA clusters; not available for individual nodes.
Total Interface Dropped Packets	Firewall	Total number of dropped packets on all interfaces. Available only at cluster level for ASA clusters; not available for individual nodes. **Note** You can view the number of per-interface dropped packets on the tabbed Interface panel presented in the detail section for the selected device.
Analysis Engine Memory (%)	IPS	Percentage of memory assigned to the Analysis Engine currently in use.
Role in Cluster	Firewall	The role of this member of an ASA load-balancing cluster: Cluster, Control Unit, or Data Unit. A cluster is managed by Security Manager as a single device with multiple nodes. Thus, each cluster is displayed in HPM as single entry, which you can expand in order to view a list of nodes.

* All of these columns are available in the All Devices and Priority Devices views.

Table Columns: VPN-related Views

You can customize the tables presented in the Monitoring pane for the VPN-related views by hiding and showing various columns of information; the columns available for display depend on the particular view.

The order of the entries in the Choose Columns to Display dialog box reflects the ordering of the columns when displayed. (However, the ordering of the rows in the following table does not necessarily reflect ordering of the columns as displayed.) See Showing and Hiding Table Columns, page 71-8 for information about opening the Choose Columns to Display dialog box.

The following table presents all available data columns for the VPN-related Monitoring views: Remote Access Users (RA), Site-to-Site Tunnels (S2S), VPN Summary, and all custom views based on these system views. Some of the listed columns are not available for specific views, as indicated.

When a device sends trap message, Cisco Security Manager captures the trap and posts it to the Site to Site Tunnels page in Health and Performance Monitoring application. Till Cisco Security Manager version 4.16, devices used to send the trap messages over IPv4 only. Beginning with Cisco Security Manager 4.17, the SNMP traps are captured using IPv6. Thus, Cisco Security Manager receives the trap messages; maps the IPv6 address of the device to the device details; displays an alert about the trap in Site to Site Tunnels page of the Health and Performance Monitoring application. The status is displayed in the Health and Performance Monitoring application after couple of refresh cycles.

Note Beginning with Security Manager version 4.9, the Health and Performance Monitoring application monitors and displays the site-to-site tunnels that have IPv6 address configured, in addition to the IPv4 based tunnels. Also the Email and Trap notifications now contain IPv6 addresses in addition to IPv4 addresses.

Table 71-2 Available Table Columns for VPN-related Views

Column Name	Available in View	Description
Receive Time	RA, S2S, VPN Summary	Poll date and time for this entry (format is: day-of-week MMM DD HH:MM:SS your-time-zone YYYY).
Firewall Name	RA, S2S, VPN Summary	Name of this device, as provided in the Security Manager inventory. Column-based Filtering, page 71-17 is available.
User Name	RA	User log-in name used to establish this session. Column-based Filtering, page 71-17 is available.
User Group Policy	RA	The name of the ASA VPN user group to which this user belongs. Column-based Filtering, page 71-17 is available.
Gateway	RA	IP address of the VPN gateway to which the user is connected. Column-based Filtering, page 71-17 is available.
Assigned IP	RA	Private IP address assigned to the remote client for this session; also known as the “inner” or “virtual” IP address.
Public IP	RA	Publicly routable IP address assigned to the client. Column-based Filtering, page 71-17 is available.

Table 71-2 Available Table Columns for VPN-related Views (continued)

Column Name	Available in View	Description
Connection Initiation Time	RA	Time and date (HH:MM:SS day-of-week MMM DD YYYY) when connection was initiated. Time is displayed in 24-hour Coordinated Universal Time (UTC) notation.
Duration	RA	Elapsed time (HH:MM:SS) between the session initiation and the most-recent device poll.
Client Version	RA	VPN client software, and version, running on the remote peer; for example, AnyConnect Windows 3.0, or Mozilla 4.0. Column-based Filtering, page 71-17 is available.
EndPoint OS	RA	Operating system in use on remote peer; for example, Windows or Windows NT. Column-based Filtering, page 71-17 is available.
Authentication Method	RA	User password, certificate, or preshared key. Column-based Filtering, page 71-17 is available.
Encryption	RA, S2S	Data encryption algorithm this session is using. Column-based Filtering, page 71-17 is available..
Tunnel Type	RA, VPN Summary (as "Type" only)	Type of tunnel or connection. These include Clientless, IPsec, and AnyConnect. Column-based Filtering, page 71-17 is available.
Throughput (Kbps)	RA, S2S	Bytes received plus bytes transmitted, in kilobits per second.
Session ID	RA	Identifier assigned to this session.
Inactive Time	RA	Amount of time this session has been inactive.
IP Address	S2S, VPN Summary	IP address of this device. Column-based Filtering, page 71-17 is available.
Local Endpoint	S2S	IP address of local tunnel interface.
Remote Endpoint	S2S	IP address of remote tunnel interface.
Local Subnet	S2S	Address of local protected subnet.
Remote Subnet	S2S	Address of remote protected subnet.
Uptime	S2S	Current duration of this tunnel.
Connection Time	S2S	Time and date (HH:MM:SS day-of-week MMM DD YYYY) when connection was initiated. Time is displayed in 24-hour Coordinated Universal Time (UTC) notation.

Table 71-2 Available Table Columns for VPN-related Views (continued)

Column Name	Available in View	Description
Status	S2S	Tunnel connection status; this will be Up or Down. An alert is issued when a tunnel goes down a specified number of times; see Alerts: Configuring, page 71-34 for more information. **Tip** You can click on a Down notification hyperlink in the Status column to view the IPSec VPN Events for that device in Event Viewer. Event Viewer will show IPSec VPN Events for the device within a time range depending on the polling interval for that device. If it is a priority device, the time range will be 5 minutes before until 5 minutes after the first down notification was received. For non-priority devices, the time range will be +/- 10 minutes instead of 5 minutes.
Health Status	VPN Summary	Current overall health of the underlying device: Critical, Warning, or Normal. Column-based Filtering, page 71-17 is available.. **Note** Overall health is defined by the most critical of any of the health metrics. For instance, if all the selected metrics on the device are normal except for one that is critical, overall device health becomes critical.
Connection Status	VPN Summary	Remote connection status; this will always be Connected. (HPM cannot present information about previous connections.) Column-based Filtering, page 71-17 is available.
Monitoring Type	VPN Summary	Types of VPN connections being monitored. Column-based Filtering, page 71-17 is available.
Active Sessions	VPN Summary	Current active sessions (S2S, IPSec RA, client-based SSL RA, and clientless SSL RA).
Peak Sessions	VPN Summary	Peak numbers of concurrent sessions (S2S, IPSec RA, client-based SSL RA, and clientless SSL RA).
Total Users	VPN Summary	Current remote user total (S2S, IPSec RA, client-based SSL RA, and clientless SSL RA).
Inactive Sessions	VPN Summary	Number of inactive sessions.
Total VPN Throughput (Kbps)	VPN Summary	Sum of all VPN traffic; that is, sum of RA and S2S throughput values, in kilobits per second. Column-based Filtering, page 71-17 is available.

Table 71-2 Available Table Columns for VPN-related Views (continued)

Column Name	Available in View	Description
ACL Name	Site-to-Site Tunnels	Beginning with version 4.9, Security Manager enables you to view the Access Control List (ACL) name that is associated with the selected Site-to-Site tunnel. This column name is selected by default. **Note** In the Health and Performance Alerts Display, Tunnel up/down Alert now also displays the ACL Name in the Description column. Similarly, Email and Trap notifications also display the ACL Name in the Description column.
Remarks	Site-to-Site Tunnels	(Optional) This column displays the remarks corresponding to the ACL Name. **Note** Alert, Email, and Trap notifications do not contain Remarks as part of the description field.

Limitation:

When Cisco Security Manager Daemon Manager service is started, the HPM application uses the latest configuration from the Configuration Archive to extract the ACL Name and Remarks associated with the Site-to-Site VPN tunnel. When the VPN tunnel is identified by HPM, it uses the extracted data to display the ACL Name and Remarks columns in the S2S view. If the VPN tunnel comes up before the data is available in HPM, the ACL Name and Remarks columns may not display the data until the next UI refresh. Similarly, the Alerts Display may not show the ACL Name in the Description column if the Alert is generated before the data is extracted by HPM. This might occur during upgrade to Security Manager version 4.9 from a previous version. If the same Alert appears in the next polling, the ACL Name is added to the Description.

Tip:

Sometimes you may notice discrepancy in the content of the Remarks column. Check if the latest configuration from the Configuration Archive contains the Remarks. If the Remarks are added or updated by Out-of-Band changes, you must perform rediscovery of the device.

Alert Table Columns

You can customize the Alerts table by hiding and showing various columns of information.

The order of the entries in the Choose Columns to Display dialog box reflects the ordering of the columns when displayed. (However, the ordering of the rows in the following table does not necessarily reflect ordering of the columns as displayed.) See Showing and Hiding Table Columns, page 71-8 for information about opening the Choose Columns to Display dialog box.

Table 71-3 Available Data Columns for the Alerts Table

Column Name	Description
Device Name (always selected)	Name of this device on which this alert was triggered, as provided in the Security Manager inventory. Column-based Filtering, page 71-17 is available.

Table 71-3 Available Data Columns for the Alerts Table (continued)

Column Name	Description
Node	The Node Name if this alert was generated by a member of an ASA load-balancing cluster Column-based Filtering, page 71-17 is available.
Device Type	Type of device: ASA or IPS. Column-based Filtering, page 71-17 is available.
Severity	Alert severity: Critical, Warning, or Normal. Column-based Filtering, page 71-17 is available.
Status	Current device status: Active or Acknowledged. Column-based Filtering, page 71-17 is available.
Description	Description of the alert. For example, “Device Health Critical” or “Device Polling: Authentication Error.”
First Seen	Date and time when this alert was first logged (day-of-week MMM DD, YYYY HH:MM:SS AM/PM). Time is based in your time zone. Column-based Filtering, page 71-17 is available.
Last Seen	Date and time when this alert was first logged (day-of-week MMM DD, YYYY HH:MM:SS AM/PM). Time is based in your time zone. Column-based Filtering, page 71-17 is available.
Notes	You can annotate an alert when you acknowledge it. Any annotations are displayed in this field. See Alerts: Acknowledging and Clearing, page 71-42 for more information.

Column-based Filtering

You can filter the various tables in HPM based on the contents of specific columns. When you apply a column filter, the table is filtered to include only those entries with the specified criteria in that column.

Note See Working with Table Columns, page 71-8 for other methods of altering table displays.

Tips

- Column filters are cumulative: for an entry to appear in the filtered table, it must meet all column filter criteria. You cannot create a set of ORed column filters.
- You can filter on the contents of most but not all columns. If a column does not have a down arrow, you cannot filter on it. For example, you cannot filter on Receive Time in All Devices view.
- The filter icon (a funnel) appears in the heading of a filtered column.
- For a description of the available columns, see Showing and Hiding Table Columns, page 71-8.

To filter a table according to a particular column parameter:

Step 1 Click the down-arrow in the heading of a column and choose one of the following from the drop-down menu:

- **All** – Choose **All** to remove or “undo” a filter from this column. The table is updated to show all entries for this parameter. For example, if you filtered the Severity column of the Alerts table to display only Critical alerts, choosing this option will re-display all Critical and Warning alerts.

- **Custom** – Choose **Custom** to open the Custom Filter dialog box where you can create a custom filter based on the information in that column. See Custom Filtering, page 71-18 for more information.
- A specific entry – The drop-down menu includes all values relevant to the column; choose one to display only that group of entries. For example, choosing **Critical** from the Severity column of the Alerts table filters the table to display only Critical alerts.

Custom Filtering

The following procedure explains how to create a custom column-based filter, one in which you are not simply selecting a value from the column's drop-down list. Refer to Column-based Filtering, page 71-17 for information about other column-based filtering options.

Step 1 Click the down-arrow in the heading of a column and choose **(Custom)** from the drop-down menu.

The Custom Filter dialog box for that column opens.

Step 2 In the Custom Filter dialog box, select the desired values. The following illustration shows a typical example of this dialog box.

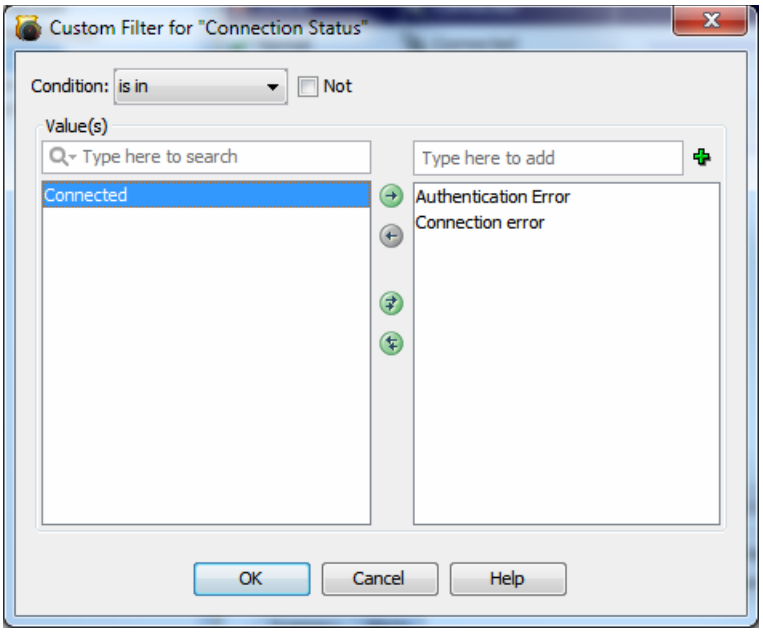

These are the controls you might find in the Custom Filter dialog box (not all controls appear for every instance):

- **Condition** – Choose the condition applied to the selected Values.

  Typically this is **is in**, meaning each of the Values you select must be "in" a column in order for that entry to be displayed in the filtered table.

- **Not** – Check this box to create a negative Condition.

  With **is in** as the chosen Condition, this would mean the selected Values cannot be in the column. In other words, the table is filtered such that entries with these Values in the column are not displayed.

- **Values** list – A few instances of the dialog box present one list of Values from which to select: simply check the desired options.

  Available and selected **Values** lists – In most cases, the dialog box presents two Values lists, as shown in the previous illustration. To select a value for the custom filter, highlight it in the left list, which contains available values for the column, and click the right arrow to add it to the list of selected values on the right. You can select multiple values.

  The items in the available Values list are determined by the values currently present in the selected column of the source table.

  If there are a lot of available values, you can search for a specific value by typing in the List Filter field above the list. For more information, see Using The List Filter Fields, page 71-19.

  You can also select, or deselect, values using the following techniques:

  - Type a Value name into the text field above the selected Values list and click the + button; the Value is added to the selected Values. This technique is useful if there is a large number of available Values, or if you want to filter on a value that is not present in the available Values list.
  - Double-click an item in either list to move it to the other list.
  - Click one of the double-arrow buttons to move all items from one list to the other, regardless of any selected values.

Step 3 Click **OK** to close the dialog box.

The table is updated to show only those entries that satisfy all currently applied filters.

Using The List Filter Fields

A List Filter field is provided above the devices and VPNs lists in the Monitoring display, above the alerts table in the Alerts display, above the device list on the VPN page of the Device Selector, and in the View Cleared Alerts window. In each case, you can use the List Filter field to quickly locate any entries in the related table that contain a specified text string.

Note The found text can be part of any data field associated with an entry. For example, as you type “license” into the Alerts List Filter field, the Alerts table is filtered to show only those alerts related to imminent license expiration. (Any matched entries are listed even if the relevant data column—in this example, Detail—is not displayed, which could cause confusion. See Showing and Hiding Table Columns, page 71-8 for more information about hiding table columns.)

Figure 71-2 Health and Performance Monitor: List Filter Field

1	Filter-parameters button.	2	Clear button.

To search for a specific text string in the devices list, the VPNs list, the Alerts table, or the View Cleared Alerts window:

- Click in the List Filter field to place the text cursor, and then begin typing.

  These are "live filter" fields. That is, as you type each character, entries that do not include your current text string are removed from the list or table. For example, suppose in an extensive list of alerts there is one with a Status of "Device Health Critical," and that none of the other alerts include any text strings containing the letters *hea*. You want to use the List Filter field to quickly locate that one alert, so you begin to enter the word "health." That alert is the only one displayed after you have typed the first three letters.

To clear a List Filter field:

- Click the clear button at the right side of the field.

  This button appears when you begin typing in the field. (You also can highlight the characters and press the Delete or Backspace key on your keyboard.)

  When you clear the List Filter field, all entries in the list are again displayed.

You can tune the filter results by specifying the information (columns) searched, by selecting case sensitivity or insensitivity, by allowing wildcards or regular expressions, and by specifying where in a returned string your characters must be located.

To change the List Filter criteria:

1. Click the filter-parameters button (magnifying glass) at the left side of the List Filter field to open the parameters menu.
2. Choose an option.

   The menu consists of four sections:

   - A list of all available information types—these entries correspond to the columns that can be displayed for that particular list or table. You can choose **All**, or alternatively you can choose individual entries.
   - **Case sensitive** and **Case insensitive** – Choose one or the other. If you choose **Case sensitive**, found text must match not only the characters you enter, but also their as-typed case.
   - **Use wildcards** and **Use regular expression** – Choose one or the other. The following wildcards are recognized:
   - ***** (asterisk) – Match zero or more characters at that location in the string.
   - **?** (question mark) – Match one character at that location in the string.

– **Match from start**, **Match exactly**, and **Match anywhere** – Choose one. **Match from start** means that the string you enter must be found at the beginning of an entry, although it can be part of a larger set of characters. **Match exactly** requires that the string you enter exactly match the entire column entry. **Match anywhere** means the string can be found anywhere within an entry, and it can be part of a larger set of characters.

3. Repeat Steps 1 and 2 to change another parameter.

Monitoring Devices

The HPM Monitoring display presents View controls, view panels, and detailed information about the currently selected device, as described in HPM Window: Monitoring Display, page 71-25.

To switch to the Monitoring screen:

- Click the **Monitoring** button below the HPM menu bar.

  (Click the **Alerts** button to return to the Alerts screen.)

Note See Managing Monitored Devices, page 71-5 for information about specifying the devices to be monitored.

This section contains the following topics:

- Managing Device Views, page 71-21
- HPM Window: Monitoring Display, page 71-25

Managing Device Views

"Views" provide the means to filter and organize the information displayed in the Monitoring pane of the HPM application. Various system views are provided—for example, All Devices, Firewall Devices, Remote Access Users Details, and so on—and you can create custom views that organize the information in other ways, such as geographic device location.

The left pane of the HPM main window displays a list of available views as shown in the following illustration.

Figure 71-3 Health and Performance Monitor: Views Pane

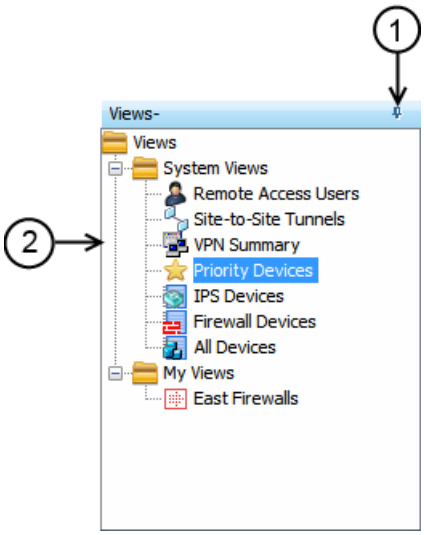

The Views pane includes the following controls:

- **(1) Push Pin button** – Click the Push Pin button to control display of the Views list. When the list is displayed as a pane of the HPM window (the pin is vertical), click the button to collapse the pane into the left edge of the window, leaving a labeled tab; the Monitoring pane is expanded to fill the HPM window.

  You can "hover" your mouse pointer over the tab to "pop out" the Views list; it remains visible as long as the pointer is over the tab or in the list area (the pin is horizontal). You also can click anywhere in the title bar—except on the pin itself—to keep the list "popped out."

  Click the pin once again to re-establish the Views list as an open pane; the Monitoring pane contracts to make room for it.

- **(2) List of views** – The list is organized into folders: System Views and My Views. Click an entry in either folder to open that view in the Monitoring pane, as described in Views: Opening and Closing, page 71-23. See Views: Custom, page 71-24 for information about creating new views in the My Views folder.

- **Right-click shortcut menu** – You can right-click any entry in the View list to access a pop-up menu of view-related commands:
  - **Edit** – Edit the name and description of the existing custom view. See Views: Custom, page 71-24.
  - **Save As** – Save the view as a new custom view. See Views: Custom, page 71-24.
  - **Delete** – Delete that custom view.
  - **Set as default view** – Use this command to designate the view that is always displayed whenever you launch the HPM application.

This section contains the following topics:

- Views: Opening and Closing, page 71-23
- Views: Tiling Horizontally or Vertically, page 71-23
- Views: Floating and Docking, page 71-24
- Views: Custom, page 71-24

Views: Opening and Closing

All available views are listed in the Views pane, on the left side of the HPM window. The Monitoring pane displays open views, with each open view presented as a separate tabbed panel. (See HPM Window: Monitoring Display, page 71-25 for more information about this window.)

Note You can detach views so they "float" in separate windows. For more information, see Views: Floating and Docking, page 71-24.

To display a new view in the Monitoring pane:

- Click the desired entry in the Views list.

  The view appears as a tabbed panel in the Monitoring pane; it is automatically selected and displayed.

To switch to another open view:

- Click the desired tab in the Monitoring pane; that view is displayed.
- Right-click any tab and choose **Next** or **Previous** to display the view to the right or left of that tabbed view.
- Click the Scroll Back and Scroll Forward buttons to the right of the tabs to display the view to the left or right of the current view.

To close a view:

- Click the close button in that tab.
- Right-click the tab and choose the **Close**.
- Right-click the tab and choose **Close Others** to close all open views except the one you right-clicked.
- Right-click any tab and choose **Close All** to close all open views.

Views: Tiling Horizontally or Vertically

Rather than displaying a single view such that it fills the Monitoring pane, you can tile two or more of the views, either horizontally or vertically, for easy comparison.

For example, if you tile two views horizontally, one view fills the upper half of the Monitoring pane, while the other fills the lower half. Similarly, tiling two views vertically fills the left-hand half of the pane with one view, with the other view filling the right half. Further, you can tile more than two views—the pane is subdivided equally for each view.

To create two horizontal or vertical tiles:

- Right-click one of the tabs and choose **New Horizontal Group** or **New Vertical Group**.

  The selected view and the other view(s) are distributed to share the Monitoring pane equally, either horizontally or vertically depending on your choice.

  Note that if there are more than two views open when you choose one of these commands, the selected view is tiled, with the remaining group of tabbed views displayed as the other tile. You can then repeat this process with the remaining tabbed views, increasing the number of visible tiles, as desired.

You can also move an existing tile to another tile:

- Right-click the tab and choose **Move to Next Tab Group** or **Move to Previous Tab Group**.

The selected view is added to the next tile (below or to the right, depending on tile orientation), or to the previous tile (above or to the left). These commands are available only if the tiled views are arranged in a manner where such movement is possible.

To change the orientation of the views, switching from horizontal to vertical tiling, or vice versa:

- Right-click any tab and choose **Change Tab Groups Orientation**.

  This command is available only when two or more tiled views are displayed.

Views: Floating and Docking

You can detach tabbed views so they "float" as separate windows, and you can "dock" floating views, returning them to the Monitoring pane as tabbed views.

To detach a view as a floating window:

- Right-click that tab and choose **Floating**.

  A standard window opens, displaying the selected view.

To move another tabbed view from the Monitoring pane to an already-open floating-view window:

- Right-click the tab and choose the window from the **Floating to** submenu.

  The right-clicked view is added to the existing window as another tabbed panel.

To return a floating view to the Monitoring pane as a tabbed panel:

- Right-click the view's tab in the window and choose **Docking**.

  That view is returned to the Monitoring pane.

Note As a standard window, you can minimize, maximize and close a floating view, as you would any other window.

Views: Custom

The Health and Performance Monitor provides seven System Views. In addition, you can create any number of custom views, each of which is based on an existing view. You also can edit and delete custom views.

The various views are presented in the Views pane of the Monitoring display, organized into two folders: System Views and My Views (the latter folder contains your custom views). The Monitoring display is described in HPM Window: Monitoring Display, page 71-25.

Follow these steps to create a new custom view:

1. In the Views list, select the view on which the new view is to be based.

   This can be a System View or an existing custom view.

2. Choose **Save As** from the File menu to open the Save View As dialog box.

   You also can right-click the selected view and choose **Save As** from the pop-up menu to open the dialog box.

3. Provide a *Name* for the new view, and optionally a *Description*.

4. Specify the devices to be monitored for this view: check and clear entries in the device-selector area of the dialog box.

5. Click **Save** to close the dialog box and add the new view to the My Views folder.

Follow these steps to edit an existing custom view:

1. Under My Views, select the view.
2. Choose **Edit** from the File menu to open the Save View As dialog box.

   You also can right-click the selected view and choose **Edit** from the pop-up menu.
3. Edit the *Name* and *Description*, as necessary.
4. Check and clear entries in the device selector to change the devices monitored for this view.
5. Click **Save** to close the dialog box.

Follow these steps to delete an existing custom view:

1. Under My Views, select the view.
2. Choose **Delete** from the File menu.

   You also can right-click the selected view and choose **Delete** from the pop-up menu.
3. Confirm that you want the view deleted.

   That view is removed from the Views list.

HPM Window: Monitoring Display

The HPM window provides two different information displays: Monitoring and Alerts. Click the Monitoring button to access the Monitoring display.

The Monitoring display consists of two primary panes: Views and Monitoring. The Views pane presents a list of available views. Click an entry in this list to open that View as a tabbed panel in the Monitoring pane.

The Monitoring pane can present multiple tabbed views, most of which display several sections. Click a tab to bring that view to the front.

Note The Remote Access Users and the Site-to-Site Tunnels views each display only a single table of information, as described in Monitoring Views: VPN, RA and S2S, page 71-30. The following descriptions focus mainly on the other available system views.

The following illustration presents the primary features of the Monitoring display and the panel sections.

Figure 71-4 Health and Performance Monitor: the Monitoring Display

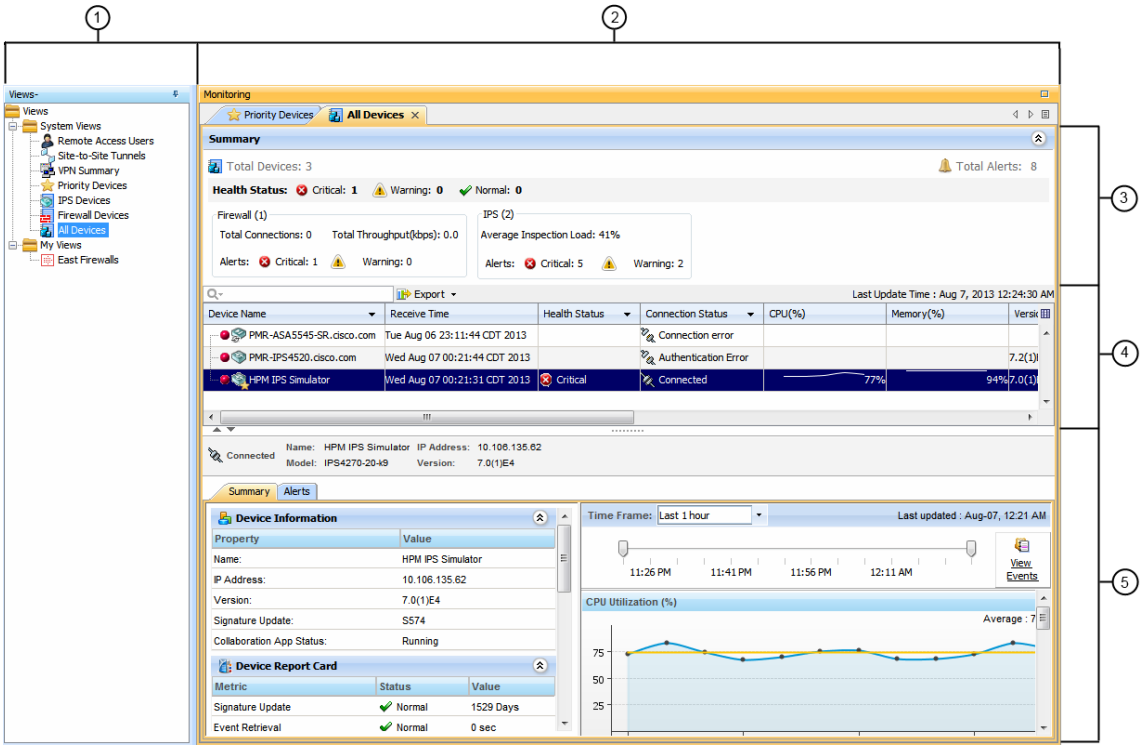

1	Views list.	**4**	Status of devices or VPNs.
2	Monitoring view controls.	**5**	Selected device details.
3	Summary of all devices.		

The Monitoring display consists of five main elements:

- **Views list (1)** – This pane lists all views available—click an entry in this list to open that view in the Monitoring pane. The views are organized into System Views, provided as part of the Health and Performance Monitor, and My Views, which are custom views you have created. See Managing Device Views, page 71-21 for information about the Views pane, and Views: Custom, page 71-24 for information about managing custom views.
- **Monitoring view controls (2)** – A labeled tab appears here for each view you open; click any tab to bring that view to the front. You also can use the Scroll Backward and Scroll Forward buttons to step backward or forward through the tabbed views. Alternately, open the Show List drop-down menu on the right and choose a label to make that the active view.
- **Summary of all devices or VPNs (3)** – Provides aggregate information for all devices or VPNs represented by this view. Expand or collapse this section by clicking the button on the right side. The device-summary section is described in greater detail in Monitoring Views: Devices or VPNs Summary, page 71-27.
- **Device-status list (4)** – All devices or VPNs included in this view are listed here; see Monitoring Views: Device or VPN Status List, page 71-27 for more information about this list. Use the List Filter field in this section to filter the list, as described in Using The List Filter Fields, page 71-19.

- **Selected device or VPN details (5)** – This section provides detailed information about the device or VPN currently highlighted in the device list. The details section is described in greater detail in Monitoring Views: Device or VPN Details, page 71-28.

This section contains the following topics:

- Monitoring Views: Devices or VPNs Summary, page 71-27
- Monitoring Views: Device or VPN Status List, page 71-27
- Monitoring Views: Device or VPN Details, page 71-28
- Monitoring Views: VPN, RA and S2S, page 71-30
- Exporting HPM Data, page 71-31

Monitoring Views: Devices or VPNs Summary

The HPM Monitoring display presents tabbed views, each of which provides detailed information about the device or VPN currently selected, as described in HPM Window: Monitoring Display, page 71-25. All device-related views (that is, all but the Remote-Access Users and Site-to-Site Tunnels views), include a Summary section, as described here.

This devices summary, or VPN Summary, which you can show and hide by clicking the button on the right side of its title bar, displays a snapshot of the aggregate Health Status and Alert Status for all the devices or VPNs relevant to the current view. For example, if you are viewing the Firewall Devices panel, the status summaries are for all monitored firewall devices only.

Monitoring Views: Device or VPN Status List

The HPM Window: Monitoring Display, page 71-25 presents detailed information about the device or VPN currently selected (in a specific device view, or the VPN Summary view, respectively). All device-related views and the VPN Summary view include a table of monitored devices or VPNs relevant to the current view.

This table displays “at-a-glance” status information for every monitored device or VPN—each is represented by an entry in this table. Note that ASA clusters are presented as expandable entries: click the + icon in front of the cluster entry to expand it and view indented entries for each cluster node.

Again, the list includes only those elements relevant to the current view. For example, the list in the Firewall Devices view does not include entries for IPS devices. The Remote-Access Users and Site-to-Site Tunnels views do not include this status display.

You can resize the table columns, you can show and hide columns, and the column headings are menus you can use to filter the table by hiding or showing devices according to chosen parameters. See Showing and Hiding Table Columns, page 71-8 for more information about these options.

When you select an entry in this list, detailed information for that device is displayed in the device-details area below the table, as described in Monitoring Views: Device or VPN Details, page 71-28.

Tip With the All Devices, Firewall Devices, IPS Devices, and Priority Devices views (and any custom device-related views), you can right-click the highlighted entry and choose **Device Manager** from the pop-up menu to open the appropriate external device manager for that device—that is, ASDM for an ASA, and IDM for an IPS sensor—where you can “drill down” into the health and performance data for that device. See Starting Device Managers, page 72-14 for more information about the device managers.

Monitoring Views: Device or VPN Details

The HPM Window: Monitoring Display, page 71-25 presents views and detailed information about the currently selected device or VPN. All device-related views and the VPN Summary view provide three or four tabbed panels of detailed information for the individual device or VPN currently selected in the device-status table above it. (The Remote-Access Users and Site-to-Site Tunnels views do not provide this details panel.)

The information presented for each type of view follows.

For the All Devices, Firewall Devices, IPS Devices, Priority Devices, and custom device-related views, the tabbed panels are:

- **Summary** - The Summary tab consists of four sections that provide information about the device and the device's status:
  - **Device Information** – This section provides a read-only listing of device-specific information such as device name, IP address, device type and model number, and so on. A read-only listing of Failover information is also presented. If an ASA cluster is selected, the Failover listing is replaced with a listing of cluster-related information.
  - **Device Report Card** - This section provides a collection of metrics that indicate the current status of the device. For more information about the metrics shown here, see Table Columns: Device-related Views, page 71-8.
  - **Interface Status** – This section provides a listing of all interfaces defined on the device, with current status information.
  - **Device Health Graphs** - This section provides a "snapshot" of device status using graphic displays for certain metrics such as CPU and memory usage. It also presents device-specific traffic information, for example, average number of connections and number of translations for firewall devices (over the most-recent polling period), and average inspection load and percentage of missed packets for IPS sensors (over the most-recent polling period).You can specify the time frame (Last 1 hour, Last 24 hours, or Last 7 days) to use for these graphs from the Time Frame list. You can focus in on a specific time frame by using the slider bar above the graphs.To view events for the selected device, click the **View Events** button. Event Viewer opens and the Event Monitoring window lists events filtered by the selected device and the time period specified by the slider bar.

    Note For IPS devices, certain health-metric thresholds must be configured separately on the individual devices—that is, outside of HPM. Therefore, it is possible for the health of an IPS device to be critical, for example, without any indication in HPM. See Alerts Configuration: IPS, page 71-35 for additional information.

- **Alerts** - The Alerts tab lists all alerts for the selected device. You can show and hide various columns of information for each alert. See Alerts and Notifications, page 71-32 for more information about alerts. For information about the fields available on this tab, see HPM Window: Alerts Display, page 71-32.

For the All Devices, Firewall Devices and Priority Devices the tabbed panels are:

- **Summary** - The Summary tab consists of four sections that provide information about the device and the device's status:

- **Device Information** – This section provides a read-only listing of device-specific information such as device name, IP address, device type and model number, and so on. A read-only listing of Failover information is also presented. If an ASA cluster is selected, the Failover listing is replaced with a listing of cluster-related information.
- **Device Report Card** - This section provides a collection of metrics that indicate the current status of the device. For more information about the metrics shown here, see Table Columns: Device-related Views, page 71-8.
- **Interface Status** – This section provides a listing of all interfaces defined on the device, with current status information.
- **Device Health Graphs** - This section provides a "snapshot" of device status using graphic displays for certain metrics such as CPU and memory usage. It also presents device-specific traffic information, for example, average number of connections and number of translations for firewall devices (over the most-recent polling period), and average inspection load and percentage of missed packets for IPS sensors (over the most-recent polling period).You can specify the time frame (Last 1 hour, Last 24 hours, or Last 7 days) to use for these graphs from the Time Frame list. You can focus in on a specific time frame by using the slider bar above the graphs.To view events for the selected device, click the **View Events** button. Event Viewer opens and the Event Monitoring window lists events filtered by the selected device and the time period specified by the slider bar.

Note For IPS devices, certain health-metric thresholds must be configured separately on the individual devices—that is, outside of HPM. Therefore, it is possible for the health of an IPS device to be critical, for example, without any indication in HPM. See Alerts Configuration: IPS, page 71-35 for additional information.

- **Alerts** - The Alerts tab lists all alerts for the selected device. You can show and hide various columns of information for each alert. See Alerts and Notifications, page 71-32 for more information about alerts. For information about the fields available on this tab, see HPM Window: Alerts Display, page 71-32.
- **Flow-offload**- The Flow-offload tab displays basic information about the offload engine, the load percentage on offload cores and information on active offloaded flows- the number of offloaded flow created, the offloaded active flows, their rewrite rules and data.
- **Flow- offload Statistics**- The Flow-offload Statistics tab displays the counts for transmitted, received and dropped packets and statistics for the virtual NIC used.

For the VPN Summary view, the tabbed panels are:

- **VPN Usage** – Several graphs presenting information such as active site-to-site tunnels, active remote-access sessions, and total throughput. This includes historical trending information for active Site-to-Site tunnels, active IPSec remote-access users, active SSL VPN clientless users, and active SSL VPN with client users.
- **Cluster Resource Usage** – Displays the usage details of cluster resources—resource names, its current, peak, and usage limits. It also provides the number of denied packets and the contexts. This feature is applicable only for Cisco Firepower 9K devices.
- **Cluster Distribution Details** – Displays the cluster mode of the VPN. If it is centralized, an error message is displayed mentioning that the destination mode is not VPN distributed. If it is distributed, it displays the respective member I and member II details. This feature is applicable only for Cisco Firepower 9K devices.

- **License Information** – A read-only listing of license information by VPN type, or IPSec and SSL license and load information, depending on your selection in the table above. For the System context of a multiple-mode device, VPN Licensing and allocation are shown; for individual contexts, VPN allocation Limits and VPN licensing usage are shown.
- **Other Details** – A listing of certificate and TrustPoint details.

See Managing Monitored Devices, page 71-5 for information about selecting devices for VPN monitoring.

Monitoring Views: VPN, RA and S2S

The HPM Monitoring display presents a variety of device- and VPN-related data views, as described in HPM Window: Monitoring Display, page 71-25. These include the Remote Access Users and Site-to-Site Tunnels views, which unlike the other views, are simply tables of current users and tunnels.

See Managing Monitored Devices, page 71-5 for information about selecting devices for VPN monitoring.

In both of these views, you can resize the table columns, you can show and hide columns, and the column headings are menus you can use to filter the table by hiding or showing entries according to chosen parameters. See Showing and Hiding Table Columns, page 71-8 for more information about these options.

The Remote Access Users view lists the remote-access users currently logged into network resources via the devices being monitored by HPM. Note that remote-access user information is updated every 20 minutes (for normal monitoring; for Priority monitoring the interval is 15 minutes), rather than the five minutes that is standard for the other views. Also, no historical or trending data is available for remote-access users.

Further, you may notice a mismatch between RA user count in the VPN Summary view and the Remote Access Users view. This is because the VPN Summary is updated at ten-minute/five-minute (normal/Priority) intervals.

Tip In the Remote Access Users view, you can right-click a user entry and choose **Log Off User** from the pop-up menu to terminate that remote-access connection.

The Site-to-Site Tunnels view provides current VPN tunnel information through all monitored devices. Note that to enable tunnel Up/Down alerts for a device or context, you must configure SNMPv3 on the device, as described in SNMP Credentials Dialog Box, page 3-48.

Tip In the Site-to-Site Tunnels view, you can click on a Down notification hyperlink in the Status column to view the IPSec VPN Events for that device in Event Viewer. Event Viewer will show IPSec VPN Events for the device within a time range depending on the polling interval for that device. If it is a priority device, the time range will be 5 minutes before until 5 minutes after the first down notification was received. For non-priority devices, the time range will be +/- 10 minutes instead of 5 minutes.

For clusters of ASA 9.0+ devices, information is shown for the control unit only, since VPN processing is not load-balanced across the nodes and is thus limited to centralized support in the cluster.

Note VPN polling occurs on a fixed time interval, so it is not possible to log status changes within that time interval. For example, if a site-to-site tunnel goes down immediately after polling and comes back up just before the next poll, that status change cannot be detected.

Exporting HPM Data

You can save a "snapshot" of the device-status information in the current View as a PDF, HTML, or CSV (comma-separated values) file.

Note Beginning with Security Manager version 4.9, the exported data in PDF, HTML, or CSV format also contain the IPv6 tunnel information.

The following steps describe exporting the current View data in either a PDF, HTML, or CSV file:

Related Topics

- HPM Window, page 71-6
- Showing and Hiding Table Columns, page 71-8

Step 1 Click the appropriate tab to display the View you want to export (that is, Priority Devices, VPN Summary, All Devices, or another).

Tip To export the data for a subset of all entries in a particular view, create a custom view that includes only the desired devices. See Views: Custom, page 71-24 for information.

Step 2 Click the down-arrow beside the Export button next to the List Filter field (above the device- or VPN-status list) and choose **As PDF**, **As HTML**, or **As CSV** from the drop-down menu.

The Export dialog box opens.

Step 3 Select the specific information to be exported by checking the appropriate columns in the dialog box.

The following topics describe the individual columns available for various views:

- Table Columns: Device-related Views, page 71-8
- Table Columns: VPN-related Views, page 71-13

Step 4 If you chose **As PDF** from the Export drop-down list, at the bottom of the Export dialog box you can choose the desired **Page Size** for the PDF file: A1, A2, A4, Letter, or Legal.

The pages of the PDF file will be the selected size, with the presented information formatted accordingly.

Step 5 If you chose **As CSV** from the Export drop-down list, Security Manager exports the information in a CSV file that you can save as required. Beginning with version 4.8, Security Manager provides the Export Trend Charts checkbox that you can select to export trend information in CSV file format. You can then chose the Time Frame from available time range of last one hour, last 24 hours, and last seven days.

Step 6 Click **Export** to close the Export dialog box.

The Save file dialog box opens.

Step 7 Provide a name for the file, and specify where it is to be saved.

The default file name is the current system time (as a long integer); you can change this to something informative. On Windows systems, the default location is My Documents; you can specify any location.

Step 8 Click **Save** to close the Save dialog box and export the selected data.

Alerts and Notifications

The Health and Performance Monitor (HPM) provides trend information, alerts, and notifications regarding the performance and health of monitored devices. You can monitor the overall health of your network—including network user and device resource utilization—by quickly scanning the status of individual devices and groups of devices.

Specific device-level trend information is available for hourly, daily and weekly intervals. Alerts are displayed prominently, with easy navigation to the relevant HPM data. You also can acknowledge and annotate individual alerts.

Note When a node from a cluster is deleted and then the cluster is rediscovered in Security Manager, the node will be removed from monitoring in HPM if currently enabled. However, any alerts generated on the node will still be shown in HPM. You must manually clear the alerts from HPM.

These alerts are based on threshold values and state-change rules that you have configured: you specify thresholds that define Critical, Warning, and Normal levels for various metrics, and you can configure rules for certain state changes such as interface failure.

Further, there are two levels of device monitoring. Initially all devices are unmonitored. However, you can designate devices to be monitored at a “normal” level, or at a “Priority” level—you define a separate set of alert definitions for each level. Priority devices are polled and reported on more frequently (five-minute intervals versus ten for “normal” devices), and failure parameters are more stringent.

You also can enable email alert notifications. If configured, an email is sent to the specified address(es) whenever an alert is generated. You can provide multiple addresses for each category of alerts (Firewall and IPS).

Note An email notification is sent the first time an alert is logged, and when the severity of an alert changes from warning to critical (but not vice-versa). No notification is issued if a device returns to the Normal state.

This section contains the following topics:

- HPM Window: Alerts Display, page 71-32
- Alerts: Configuring, page 71-34
- Alerts: Viewing, page 71-41
- Alerts: History, page 71-43

HPM Window: Alerts Display

The HPM window provides two different information displays: Monitoring and Alerts. Click the Alerts button to access the Alerts display.

Note A device-specific view of alert data is available on the Alerts tab when viewing details for a specific device (see Monitoring Views: Device or VPN Details, page 71-28). With a few exceptions, you can perform many of the same functions from the device-specific alert view as you can from the primary Alerts display.

The following illustration presents the primary features of the Alerts display.

Related Topics

- Alerts: Configuring, page 71-34

Figure 71-5 Health and Performance Monitor: Alerts Display

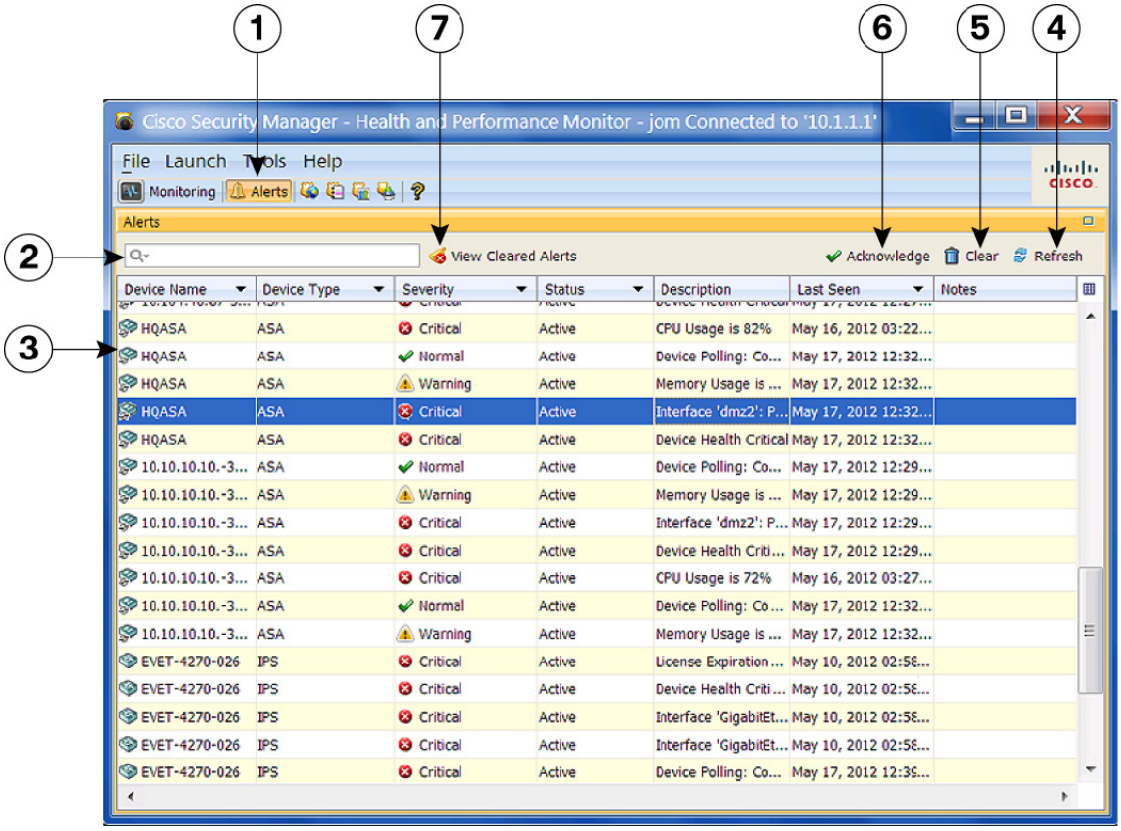

1	Alerts button.	**5**	Clear button.
2	List Filter field.	**6**	Acknowledge button.
3	Alerts table.	**7**	View Cleared Alerts button.
4	Refresh button.		

The Alerts display consists of seven main elements:

Note With the exception of the List Filter field and the View Cleared Alerts button, these same elements are available to you on the Alerts tab when viewing details for a specific device (see Monitoring Views: Device or VPN Details, page 71-28).

- **Alerts button (1)** – The HPM window displays either Monitoring information for devices and VPNs, or a table of alerts generated by monitored devices. Click the Alerts button to view the alerts table.
- **List Filter field (2)** – You can use this field to filter the alerts displayed in the table; only those alerts containing the specified text are listed. Refer to Using The List Filter Fields, page 71-19 for more information.
- **Alerts table (3)** – This table lists all alerts for all currently monitored devices. The alerts displayed can be filtered using the List Filter field. You also can show and hide various columns of information for each alert. See Alerts and Notifications, page 71-32 for more information.
- **Refresh button (4)** – Click this button to update all alerts ahead of the normal polling cycles.
- **Clear button (5)** – When one or more alerts are selected, you can click this button to open the Clear dialog box. Click the Clear button in the dialog box to close it and clear the highlighted alerts from the table.

  Note See Alerts: Acknowledging and Clearing, page 71-42 for additional information about clearing and acknowledging alerts.

- **Acknowledge button (6)** – When one or more alerts are selected, you can click this button to open the Acknowledge dialog box. If desired, you can enter a note that will be applied to the selected alerts. Click the Acknowledge button to close the dialog box and mark all highlighted alerts as acknowledged.

Tip You can add a note to any previously acknowledged alert. Click the Note field for that alert to open the Enter Notes dialog box. This is the only method of accessing the Enter Notes dialog box.

- **View Cleared Alerts button (7)** – Click this button to open the View Cleared Alerts window where you can access and view previously cleared alerts; you specify a set of devices and a time range. See Alerts: History, page 71-43 for more information about using this window.

Alerts: Configuring

The alerts and email notifications provided by HPM are based on threshold values and state-change rules that you configure in the Alerts Configuration dialog box.

The Alerts Configuration dialog box consists of three tabbed panels: **IPS** for IPS sensor-related alerts, **FW** for firewall-related alerts, and VPN for tunnel-status alerts. Each panel presents groups of options in sections—use the expand/collapse button to show or hide a particular section.

Note You can enable and disable a particular alert without expanding that section; simply check or clear the box preceding the section heading—the current settings are used and retained.

There are two levels of device monitoring: normal or “standard” priority and “active” priority. Active priority devices are polled and reported on more frequently, and failure parameters are more stringent. You can designate up to 10% of all monitored devices for Priority monitoring. See Managing Monitored Devices, page 71-5 for more information about device selection.

Follow these steps to configure alert reporting and notifications for both Standard and Priority devices:

Step 1 Choose Alert Configuration from the Tools menu to open the Alerts Configuration dialog box.

Step 2 On the IPS panel, configure IPS-related alerts—if necessary, click the IPS tab to display the panel.

1. To enable email Notifications when IPS alerts are generated, enter one or more valid addresses in the Email Addresses field; separate multiple addresses with commas.
2. Use the checkboxes in the section headings to enable and disable specific alerts. Expand a section to update those alert definitions. The IPS parameters are described in Alerts Configuration: IPS, page 71-35.

Note An email notification is sent the first time an alert is logged, and when the severity of an alert changes from warning to critical (but not vice-versa). No notification is issued if a device returns to the Normal state.

Step 3 On the FW panel, configure firewall-related alerts—click the FW tab to display the panel.

1. To enable email Notifications when firewall alerts are generated, enter one or more valid addresses in the Email Addresses field; separate multiple addresses with commas.
2. Use the checkboxes in the section headings to enable and disable specific alerts. Expand a section to update those alert definitions. The FW parameters are described in Alerts Configuration: Firewall, page 71-37.

Step 4 On the VPN panel, configure tunnel-status alerts—click the VPN tab to display the panel.

1. To enable email Notifications when tunnel-down alerts are generated, enter one or more valid addresses in the Email Addresses field; separate multiple addresses with commas.
2. Use the checkbox in the section heading to enable and disable tunnel-status alerts. Expand the section to update those alert definitions. The VPN parameters are described in Alerts Configuration: VPN, page 71-39.

Note To enable these tunnel-status alerts for a device or context, you must first configure SNMP on the device, as described in Configuring SNMP for S2S Polling, page 71-40.

Step 5 Click **Save** to save your changes and close the dialog box.

Alerts Configuration: IPS

Note From version 4.17, Cisco Security Manager does not support IPS devices.

The alerts and status information collected from monitored IPS devices are configured on the IPS panel of the Alerts Configuration dialog box. Refer to Alerts: Configuring, page 71-34 for information about opening the dialog box, accessing the IPS panel, and providing email addresses for IPS-related Notifications.

The IPS-alert configuration parameters are grouped into sections that can be expanded and collapsed. Each section includes a checkbox next to its heading; use this checkbox to enable or disable that alert. When expanded, each section provides access to the settings used to define the alert.

The IPS alert and status configuration parameters are described in the following table. Each parameter can be configured separately for Priority Devices and Standard Devices. (Specifying devices for priority and standard monitoring is described in Managing Monitored Devices, page 71-5.)

Note Some of the following alert settings require specific related parameters to be configured on the monitored IPS sensors themselves. For example, if **license-expiration-policy** (**health-monitor** command) is not enabled on a particular sensor, license-expiration messages are not generated by that sensor and therefore no occurrences are tallied for it by HPM.

Table 71-4 IPS Alerts Configuration

Setting	Description
Collaboration App Status	Errors generated by the Collaboration App application are tallied. Alerts and Notifications are generated when the number of errors tallied reaches the specified Occurrences value.
Sensor App Status	Errors generated by the Sensor App application are tallied. Alerts and Notifications are generated when the number of events reaches the specified Occurrences value.
Bypass Mode	Any time bypass mode is triggered, one Occurrence is tallied for this setting. Alerts and Notifications are generated when the number of Occurrences reaches the value specified.
Interface Status	The status of each enabled interface is polled periodically. Each “down” result for any given interface is tallied as one Occurrence for that interface. Alerts and Notifications are generated when the number of Occurrences reaches the value specified.
License Expiration	A license-expiration threshold can be configured on each IPS sensor, and whenever this threshold is crossed, a status message is issued.
Memory Usage	A memory-usage threshold can be configured on each IPS sensor, and whenever this threshold is exceeded, a status message is issued. An Occurrence is tallied for each memory-usage message. Alerts and Notifications are generated when the number of Occurrences reaches the value specified here.
Missed Packets	A missed-packets threshold can be configured on each IPS sensor, and whenever this threshold is exceeded, a status message is issued. An Occurrence is tallied for each missed-packets message. Alerts and Notifications are generated when the number of Occurrences reaches the value specified here.

Table 71-4 IPS Alerts Configuration (continued)

Setting	Description
Inspection Load	A traffic inspection-load threshold can be configured on each IPS sensor, and whenever this threshold is exceeded, a status message is issued. An Occurrence is tallied for each load-exceeded message. Alerts and Notifications are generated when the number of Occurrences reaches the value specified.

Alerts Configuration: Firewall

The alerts and status information collected from monitored firewall devices are configured on the **FW** panel of the Alerts Configuration dialog box. Refer to Alerts: Configuring, page 71-34 for information about opening the dialog box, accessing the FW panel, expanding and collapsing sections, and providing email addresses for FW-related Notifications.

The firewall-alert configuration parameters are grouped into sections that can be expanded and collapsed. Each section includes a checkbox next to its heading; use this checkbox to enable or disable that alert. When expanded, each section provides access to the settings used to define the alert.

Some section headings also include **Consider for Device Health** checkboxes. Checking one of these boxes means that particular information is considered when determining overall health of each device.

The FW alert and status configuration parameters are described in the following table.

Table 71-5 Firewall Alerts Configuration

Setting	Description
Failover Peer Status	The status of the link to the device's failover peer is polled periodically. Each failed contact attempt is tallied as one Occurrence. Alerts and notifications are generated when the number of occurrences reaches the values specified here. For Priority devices and for Standard devices: choose **Critical** or **Warning** to specify the type of alert generated, and then specify the number of occurrences necessary to trigger the alert.
Interface Status	The status of each enabled interface is polled periodically. Each "down" result for any given interface is tallied as one Occurrence for that interface. This monitoring is per stand-alone device, and per node of an ASA cluster. Alerts and notifications are generated when the number of occurrences reaches the values specified here. For Priority devices and for Standard devices: choose **Critical** or **Warning** to specify the type of alert generated, and then specify the number of occurrences necessary to trigger the alert. **Note** Check **Consider for Device Health** in the header to include these data in device-health calculations.

Table 71-5 Firewall Alerts Configuration (continued)

Setting	Description
Master Changed	An Occurrence is tallied each time the device designated as control unit of an ASA cluster changes. Alerts and notifications are generated when the number of occurrences reaches the values specified here. For Priority devices and for Standard devices: choose **Critical** or **Warning** to specify the type of alert generated, and then specify the number of occurrences necessary to trigger the alert.
Cluster Node Status	An Occurrence is tallied each time the Connection Status of an ASA cluster node changes (comes up or goes down). Alerts and notifications are generated when the number of occurrences reaches the values specified here. For Priority devices and for Standard devices: choose **Critical** or **Warning** to specify the type of alert generated, and then specify the number of occurrences necessary to trigger the alert.
CPU Usage	An Occurrence is tallied each time CPU usage exceeds the specified Threshold percentage. This is per stand-alone device; per node of a single-context cluster; and per node for the system context only in a multi-context cluster. Alerts and notifications are generated when the number of occurrences reaches the values specified here. **Note** Check **Consider for Device Health** in the header to include these data in device-health calculations. For Priority devices and for Standard devices, you can enable either or both **Critical** and **Warning** CPU Usage alerts: 1. Check the appropriate box to enable the **Threshold** and **Occurrence** fields. 2. Specify a **Threshold** percentage by clicking the up or down arrows, or by highlighting the existing value and typing a number. 3. In the **Occurrence** field, specify the number of times the specified Threshold must be exceeded before the critical or warning alert is issued.

Table 71-5 Firewall Alerts Configuration (continued)

Setting	Description
Memory Usage	An Occurrence is tallied each time memory usage exceeds the specified Threshold percentage. This is per stand-alone device; per node of a single-context cluster; and per node for the system context only in a multi-context cluster. Alerts and notifications are generated when the number of occurrences reaches the values specified here. **Note** Check **Consider for Device Health** in the header to include these data in device-health calculations. For Priority devices and for Standard devices, you can enable either or both **Critical** and **Warning** Memory Usage alerts: 1. Check the appropriate box to enable the **Threshold** and **Occurrence** fields. 2. Specify a **Threshold** percentage by clicking the up or down arrows, or by highlighting the existing value and typing a number. 3. In the **Occurrence** field, specify the number of times the specified Threshold must be exceeded before the critical or warning alert is issued.

Alerts Configuration: VPN

The generation of alerts for site-to-site (S2S) tunnels on monitored devices and contexts is enabled and configured on the **VPN** panel of the Alerts Configuration dialog box. Refer to Alerts: Configuring, page 71-34 for information about opening the dialog box, accessing the VPN panel, and providing email addresses for VPN-related Notifications.

Tip When VPN alerts are enabled, HPM polls the monitored devices and contexts at normal and Priority intervals (ten and five minutes, respectively), according to your normal/Priority designations. You also can enable SNMP monitoring which updates HPM tunnel status immediately upon processing the traps. See Configuring SNMP for S2S Polling, page 71-40 for more about enabling SNMP processing for HPM.

The tunnel-status configuration parameters are grouped into a section that can be expanded and collapsed. When expanded, you have access to the alert settings. The checkbox next to the heading is used to enable or disable the alert.

The VPN alert parameters are described in the following table.

Table 71-6 VPN Alerts Configuration

Setting	Description
Tunnel Status	The status of each monitored S2S tunnel is updated whenever it comes up or goes down, based on periodic polling or SNMP trap processing. Each "down" result for any given tunnel is tallied as one Occurrence. An alert is generated when the number of occurrences reaches the values specified here. For Priority Devices and for Standard Devices, you can separately configure both Critical and Warning tunnel-down alerts: choose **Critical** or **Warning** to specify the type of alert generated, and then in the **Occurrence** field, specify the number of times a tunnel is down when polled before the critical or warning alert is issued.

Configuring SNMP for S2S Polling

The Health and Performance Monitor (HPM) application uses SNMP to poll site-to-site (S2S) VPN tunnels for up/down status updates. The generation of alerts for site-to-site (S2S) tunnels on monitored devices and contexts is configured on the **VPN** panel of the HPM Alerts Configuration dialog box. Refer to Alerts: Configuring, page 71-34 for information about opening the dialog box, accessing the VPN panel, and providing email addresses for VPN-related Notifications.

Configuring SNMP in Security Manager to provide S2S polling is outlined here. The basic steps are:

1. Enable and configure SNMP on the SNMP Page, page 49-17 for the device or individual context; specifically: check Enable SNMP Servers and provide and confirm the Read Community String.
2. In the SNMP Trap Configuration Dialog Box, page 49-19, check **IPSEC Start** and **IPSEC Stop** on the Other panel.
3. In the Add/Edit SNMP Host Access Entry Dialog Box, page 49-22, provide Interface Name, IP Address, Community String (and Confirm it), and choose the SNMP Version (1 or 2c).

   Versions 1, 2c and 3 are supported for S2S polling, but version 3 must be configured separately, as described in the next section.
4. Configure SNMP credentials for the device or individual context in the SNMP Credentials Dialog Box, page 3-48.

   For versions 1 and 2c, provide and confirm the RO Community String.

   For version 3, Security Manager supports three modes; which to use is determined from your input:

   - noauthnopriv (no authentication, no privacy) – User name is mandatory, others are optional.
   - authnopriv (authentication, no privacy) – User name, Password, Auth Algorithm, and Engine ID are required.
   - authpriv (authentication and privacy) – User name, Password, Auth Algorithm, Privacy Password, Privacy Algorithm, and Engine ID are required.

   Again, configuration of SNMP v3 is performed separately, as described in the next section.

Configuring SNMP v3 for Security Manager Device

You cannot configure SNMP v3 directly in Security Manager; you must use CLI commands or set up a FlexConfig. The steps are:

1. Configure an SNMP server group.

```
snmp-server group group-name v3 [auth | noauth | priv]
```

The `auth` keyword enables packet authentication. The `noauth` keyword indicates no packet authentication or encryption is being used. The `priv` keyword enables packet encryption and authentication. There are no default values for the `auth` or `priv` keywords.

2. Define a new SNMP user.

```
snmp-server user username group-name{v3 [encrypted]
[auth {md5 | sha]} auth-password
[priv [des | 3des | aes] [128 | 192 | 256] priv-password]
```

The `v3` keyword specifies that the SNMP Version 3 security model is used, and enables the use of the `encrypted`, `priv`, and `auth` keywords. The `encrypted` keyword indicates the password is in encrypted format. Encrypted passwords must be in hexadecimal format.

The `auth` keyword specifies which authentication level (`md5` or `sha`) is used.

The `priv` keyword specifies the encryption level. There are no default values for the `auth` or `priv` keywords.

For the encryption algorithm, you can specify either `des`, `3des`, or `aes`. You can also specify which version of the AES encryption algorithm to use: `128`, `192`, or `256`. The `auth-password` specifies the authentication user password. The `priv-password` specifies the encryption user password.

3. Specify the recipient of SNMP notifications.

```
snmp-server host interface {hostname | ip_address} [version 3 username]
```

Indicates the interface from which traps are sent. Identifies the name and IP address of the NMS or SNMP manager that can connect to the device.

Related Topics

- Configuring SNMP, page 49-14

Alerts: Viewing

All alerts generated for monitored devices are displayed as a table in an alternate screen of the HPM window. The Alerts table is updated automatically as devices are polled for status information. You can also click the Refresh button, above the table on the right side, to update the table.

These alerts are based on the threshold values and state-change rules you have configured. See Alerts: Configuring, page 71-34 for more information.

Note See Managing Monitored Devices, page 71-5 for information about specifying the devices to be monitored.

To switch to the Alerts screen:

- Click the **Alerts** button below the HPM menu bar.

  (Click the **Monitoring** button to return to the Monitoring screen.)

Note You can also view alerts that apply to a specific device from the Alerts tab when viewing details for that device (see Monitoring Views: Device or VPN Details, page 71-28).

The Alerts listing is a basic table, consisting of rows and columns, with each row representing one alert from a given device. Each column provides specific information about that alert: device name, alert severity, time recorded, and so on. (See HPM Window: Alerts Display, page 71-32 for more about the Alerts screen.)

Note The column headings are menus that you can use to filter the table by hiding or showing alerts according to chosen parameters. For example, you might choose to display alerts for only a particular device, and then choose only critical alerts for that device. See Working with Table Columns, page 71-8 for more information.

Tip You can click on the hyperlink in the Description column for tunnel up/down alerts to view the IPSec VPN Events for that device in Event Viewer. Event Viewer will show IPSec VPN Events for the device within a time range depending on the polling interval for that device. If it is a priority device, the time range will be 5 minutes before until 5 minutes after the first up/down notification was received. For non-priority devices, the time range will be +/- 10 minutes instead of 5 minutes.

In addition to scrolling the Alerts table, you can view sets of specific alerts:

- Use the List Filter field above this table to filter the list. See Using The List Filter Fields, page 71-19 for more information.
- Use the View Cleared Alerts window to view previously cleared alerts for a selected set of devices over a specified time range. See Alerts: History, page 71-43 for more information.

You also can acknowledge alerts, clear alerts, and edit alert notes:

- You can acknowledge an alert, or clear it, as described in Alerts: Acknowledging and Clearing, page 71-42.
- To add to an existing alert note, click Notes field for that entry in the table to open the Enter Notes dialog box—used to view and add notes to an alert. Available only when a single alert with an existing note is selected in the table.

Alerts: Acknowledging and Clearing

All alerts generated for monitored devices are displayed in the Alerts table, as described in Alerts: Viewing, page 71-41. You can add notes to individual alerts, and you can acknowledge or clear alerts individually or in groups.

To select an alert, click that entry in the Alerts table. You can Shift-click another alert to select the group between the two, and you can Ctrl-click various rows to select multiple non-contiguous alerts.

When an alert is selected in the table, you can:

- Click the **Acknowledge** button to open the Acknowledge Alert dialog box, used to add a note to, and then mark the selected alert(s) as acknowledged. You can acknowledge multiple alerts at one time.

  Enter text in the Notes field in this dialog box (this is optional), and then click **OK**. The dialog box closes and the alerts are marked as acknowledged with a timestamp displayed in the Notes column.
- Click the **Clear** button to open the Clear Alert dialog box, used to add a note to, and then remove the selected entries from the Alerts table.

  Enter text in the Notes field in this dialog box (this is optional), and then click **OK**. The dialog box closes and the selected alerts are removed from the Alerts table.

Note Alerts can be cleared automatically by HPM if you change the relevant threshold(s). Like alerts you have cleared, these alerts can be viewed in the View Cleared Alerts window (see Alerts: History, page 71-43).

Notes and other information for cleared alerts are saved in an Alerts database for 30 days.

Alerts: History

All alerts generated for monitored devices are displayed as a table in the HPM window. You can filter the table by any visible column parameter, as described in Alerts: Viewing, page 71-41.

You also can use the View Cleared Alerts window to access and view previously cleared alerts; you specify a set of devices and a time range. (Clearing alerts is described in Alerts: Acknowledging and Clearing, page 71-42.)

Note Notes and other information for cleared alerts is maintained in an Alerts database for 30 days—you cannot access alerts more than 30 days old.

Follow these steps to open and use the View Cleared Alerts window:

1. In the Alerts screen, click the View Cleared Alerts button next to the List Filter field to open the View Cleared Alerts window. (See Alerts: Viewing, page 71-41 for more information about accessing the Alerts screen of the HPM window.)
2. Specify the alert View Settings; these define the set of alerts you wish to view:
   - Specify the devices of interest; **All** devices are selected by default. To select a particular set of devices:
     a. Click the **Select** button to open the Select Devices dialog box.
     b. Select the desired device(s); deselect any devices you wish to exclude.
     c. Click **OK** to close the Select Devices dialog box.
   - Specify the types of Alerts to display: select or deselect **Critical**, **Warning** and **Normal**.
   - Define the desired **Time Range** by choosing a From date and time, and a To date and time. All alerts with a First Seen time within this range will be displayed.

     From and To each present a standard drop-down calendar used to select a month and day.

     Use the time field below each calendar to specify the precise start or end time, respectively. Highlight a digit and click the up or down arrow, or simply type the desired number. You can also click the **Now** button to specify the present moment.
3. Click the **Search** button to display the defined set of alerts.

Note that the View Cleared Alerts window provides a List Filter field that you can use to filter the cleared-alerts display. Using this field is described in Using The List Filter Fields, page 71-19.

Refer to Working with Table Columns, page 71-8 for other methods of filtering this table.

SNMP Trap Forwarding Notification

In 4.6 and earlier versions of Security Manager, e-mail notifications were sent to users when Health and Performance Monitor alerts were generated for ASA, IPS, and VPN.

This framework has been enhanced for Security Manager 4.7 to send SNMP trap notifications in addition to e-mail notifications. Security Manager 4.7 converts the alerts to traps and sends them to the centralized SNMP trap server. SNMP v1, v2c, and v3 are supported. A trap is generated the first time an alert is seen and again if the severity increases, resulting in a trap being generated a maximum of 2 times, as is done with e-mail notification.

SNMP trap forwarding notification has the following pre-requisites:

1. An SNMP trap receiver (server) is available. More than one server can be used for a given Security Manager installation.
2. An ASA device is available.
3. An IPS sensor running IPS 7.0.x or later is available.
4. Health and Performance Monitor is enabled.

Tip To verify that Health and Performance Monitor is enabled, navigate to Configuration Manager > Tools > Security Manager Administration... > Health and Performance Monitor.

5. Normal or priority monitoring for the ASA device and the IPS sensor Device is enabled in Health and Performance Monitor.
6. The alert settings for firewall, IPS, and VPN are enabled.

Tip To enable the alert settings for firewall, IPS, and VPN, navigate to Health and Performance Monitor > Tools > Alert Configuration.

MIB Documentation

The information in this section documents which MIB Security Manager uses to send the trap notification and what OID the user has to look for to get the particular alert information.

For the SNMP trap, Security Manager uses "CISCO-DEVICE-EXCEPTION-REPORTING-MIB"

The following list contains the OID details and the information that it contains:

- iso.org.dod.internet.mgmt.mib-2.system.sysUpTime.0 (shows how long the HPM server has been up and running) – calculated using the formula, System Up time = Current Security Manager server time - HPM Service start up time.
- snmpTrapOID (1.3.6.1.4.1.9.9.224.2.0.1)
- .1.3.6.1.4.1.9.9.224.1.1.5.1.2 (lists the alert rule name like memory usage)
- .1.3.6.1.4.1.9.9.224.1.1.5.1.3 (constant value of 1, specifies IP address type)
- .1.3.6.1.4.1.9.9.224.1.1.5.1.4 (device display name – device type and if any cluster node, then its name)
- .1.3.6.1.4.1.9.9.224.1.1.5.1.5 (severity of the alert)
- .1.3.6.1.4.1.9.9.224.1.1.5.1.6 (time stamp of the alert) – Current Security Manager server time - Alert first seen time
- .1.3.6.1.4.1.9.9.224.1.1.5.1.7 (maximum 1024 char string describing the alert)

- .1.3.6.1.4.1.9.9.224.1.1.5.1.8 (Security Manager Server name)

This section contains the following topics:

- SNMP Trap Entries Dialog Box, page 71-45
- Add/Edit/Copy SNMP Trap Entries Dialog Box, page 71-46

SNMP Trap Entries Dialog Box

Use the SNMP Trap Entries dialog box as your launching point for SNMP trap forwarding notifications.

Navigation Path

In Health and Performance Monitor, select **SNMP Trap Configuration** from the Tools menu. The SNMP Trap Entries dialog box contains the following areas:

- The Settings table, which displays the traps that are currently configured.
- Add, Edit, and other options for working with SNMP trap entries.

Field Reference

Table 71-7 The Settings table and other options in the SNMP Trap Entries dialog box

Field	Description
Forward Trap To table	
Status column	Enabled or Disabled. A maximum of 5 trap forwarding hosts can be enabled at any given time.
IP/Host column	IP address or hostname of the centralized SNMP trap server. Neither the local host nor the Security Manager server is allowed to be used as the SNMP server, in order to avoid problems with performance.
PORT column	The port used by the centralized SNMP trap server
SNMP Version column	v1, v2c, or v3. The default value is v2c.
Username column	The username for authentication to the SNMP server
Authentication Algorithm column	MD5 or SHA
Encryption Algorithm column	DES, 3DES, AES128, AES192, and AES256
Add	Used to add a new configuration. The Add button opens the Add Trap Settings dialog box.
Edit	Used to edit an existing configuration. The edit button opens the Edit Trap Settings dialog box.
Copy Settings	Used to copy all the settings in an existing configuration. The Copy Settings button opens the Copy SNMP Trap Settings dialog box.
Delete	Used to delete an existing configuration
Enable	Used to enable an existing configuration
Disable	Used to disable an existing configuration

Add/Edit/Copy SNMP Trap Entries Dialog Box

Use the Add/Edit/Copy SNMP Trap Entries dialog box to add, edit, and otherwise work with and configure SNMP traps.

Navigation Path

In Health and Performance Monitor, select **SNMP Trap Configuration** from the Tools menu. Then select **Add**, **Edit**, or **Copy Settings**.

The Add/Edit/Copy SNMP Trap Entries dialog box contains the following areas:

- The IP/Host and Port area
- The Trap Settings area for FW Alerts, IPS Alerts, and VPN Alerts
- The Trap Settings area for SNMP Options

Field Reference

Table 71-8 The Trap Settings area and other options in the Add/Edit/Copy SNMP Trap Entries dialog box

Field	Description
Trap Settings	Used to select all or only selected alerts for FW, IPS, and VPN, described in the following topics: • Alerts Configuration: Firewall, page 71-37 • Alerts Configuration: IPS, page 71-35 • Alerts Configuration: VPN, page 71-39
SNMP Options	
RO Community String (SNMP version v1 and v2c only)	The password used for authentication in SNMP version v1 or v2c.
Group Type (SNMP version v3 only)	NOAUTH, AUTH, or PRIV.
Engine ID (SNMP version v3 only)	The SNMPEngineID identifier used for authentication in v3.
User Name (SNMP version v3 only)	The username for authentication to the SNMP server.
Authentication Password (SNMP version v3 only)	The password for authentication to the SNMP server.
Authentication Protocol (SNMP version v3 only)	MD5 or SHA.
Encryption Password (SNMP version v3 only)	The password for MD5 or SHA encryption.
Encryption Protocol (SNMP version v3 only)	DES, 3DES, AES128, AES192, and AES256.

NOTE: In order to use AES192, AES256, or 3DES, you must follow these steps:

1. Download the unlimited strength cryptography policy .jar files from http://www.oracle.com/technetwork/ > Downloads > Java SE > Java Cryptography Extension (JCE) Unlimited Strength Jurisdiction Policy Files for JDK/JRE 7. (Click the download button to download the files by accepting the license agreement.)
2. Replace local_policy.jar and US_export_policy.jar on your Security Manager server in the folder CSCOpx\MDC\vms\jre\lib\security.
3. Restart your Security Manager server.

Using External Monitoring, Troubleshooting, and Diagnostic Tools

A high degree of network availability is a requirement for large enterprises and service providers. Network managers face various challenges in maintaining network availability, including unscheduled down time, lack of expertise, insufficient tools, complex technologies, business consolidation, and competing markets. Network monitoring, problem diagnosis, and troubleshooting are essential to meeting and overcoming these challenges.

Monitoring involves the study of network activity and device status to identify anomalous events and behaviors. Quickly diagnosing and correcting network and system faults such as outages and degradations increase service availability, and thus tools to isolate, analyze and correct problems are essential.

The main Security Manager tools for monitoring device events are the Health and Performance Monitor (see Chapter 71, "Health and Performance Monitoring") and the Event Viewer (see Chapter 69, "Viewing Events").

In addition to Health and Performance Monitor and Event Viewer, the following topics describe additional monitoring, troubleshooting and diagnostic tools that are available with Security Manager:

- Dashboard Overview, page 72-1
- CSM Mobile, page 72-11
- Viewing Inventory Status, page 72-12
- Starting Device Managers, page 72-14
- Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20
- Analyzing an ASA or PIX Configuration Using Packet Tracer, page 72-23
- Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools, page 72-26
- Using the Packet Capture Wizard, page 72-30
- IP Intelligence, page 72-34
- Integrating CS-MARS and Security Manager, page 72-36

Dashboard Overview

Beginning with Version 4.5, the Security Manager client has a new launch point—a configurable dashboard, for which this topic presents an overview.

The dashboard is one of the six client applications that you can select as your default client application when you start the Security Manager client. (The others are Configuration Manager, Event Viewer, Report Manager, Health and Performance Manager, and Image Manager; there is also an application designed for mobile devices called CSM Mobile.) The dashboard is a convenient way for you to accomplish tasks that are found in several other areas of Security Manager, such as the IPS Health Monitor page, Report Manager, Health and Performance Monitor, and IP Intelligence Settings.

The dashboard contains the widgets shown in the following table, categorized by whether they are for use with IPS, firewalls, or both. (Not all of these widgets are shown by default). In addition to the original dashboard, you can create new, additional dashboards, which are displayed as tabs. You can customize all dashboards, both the original dashboard and any new, additional dashboards that you create. To customize a dashboard, you can drag and drop widgets from the list of available widgets into any dashboard.

Table 72-1 Widgets for IPS, Firewalls, and Both

Widgets for IPS	• IPS Inspection Load Trends • Top 10 Reports for IPS Attackers, Victims, and Signatures • IPS Missed PacketTrends • IPS License • IPS Update Packages • IPS Sensors Out of Date
Widgets for Firewalls	• Top 10 Reports for Firewall Sources, Destinations, and Services • Top 10 Reports for Botnet Malware Sites, Ports, and Hosts • Firewall CPU Usage Trends
Widgets for Both IPS and Firewalls	• Device Health Summary • Memory Usage Trends • Deployment • IP Intelligence

The way in which you use the dashboard and its widgets depends upon your goals in using Security Manager. For example, you can use the following four widgets to observe device health trends:

- IPS Inspection Load Trends
- IPS Missed Packet Trends
- Memory Usage Trends
- Firewall CPU Usage Trends

Individual widgets are described in the following table. One of the key widgets is the Device Health Summary widget. One reason it is important is that it provides the same information accessible through CSM Mobile, which is designed specifically for mobile devices. For more information about CSM Mobile, see CSM Mobile, page 72-11. For information on enabling or disabling CSM Mobile, see CSM Mobile Page, page 11-9.

Table 72-2 Description of Individual WidgetsDashboard widgetswidgets for IPS in Dashboardwidgets for firewall in dashboard

IPS Inspection Load Trends	A measure of the IPS inspection load trends. The inspection load trend data will appear in this widget only when an IPS device issues an alert because of inspection load, and the data will disappear when the alert is cleared. Indicates how much traffic inspection capacity the sensor is using. 0 indicates that there is no traffic backup, and 100 indicates that the buffers are completely backed up. Inspection load trends are affected by the following things: • Rate of traffic that needs inspection • Type of traffic being inspected • Number of active connections being inspected • Rate of new connections per second • Rate of attacks being detected • Signatures active on the sensor • Custom signatures created on the sensor You can set monitoring parameters on the IPS Health Monitor page at [IPS device in Device View] Platform > Device Admin > Health Monitor.

Table 72-2 Description of Individual WidgetsDashboard widgetswidgets for IPS in Dashboardwidgets for firewall in dashboard (continued)

Top 10 Reports for IPS Attackers, Victims, and Signatures	Pre-defined system reports that you can use to analyze top attackers, victims, and signatures for IPS alerts in your network. **Clickable Link**—In the Top Attackers widget, the IP address is an active hyperlink; click it to display IP intelligence. For details on IP intelligence in Security Manager, refer to IP Intelligence, page 72-34. **Clickable Link**—In the Top Signatures widget, the Signature ID is an active hyperlink; click it to display signature information. To use these reports, use Report Manager (Launch > Report Manager...). To cross-launch Event Viewer from one of these top ten reports, select a particular attacker, victim, or signature, and click the number of occurrences. The number of occurrences is listed for the last 24 hours by default; you can change it to the last hour if desired. **Note** When you cross-launch Event Viewer, the event query time in Event Viewer will be shown as the last 10 minutes despite its being the last 24 hours or the last hour in the Summary Dashboard. You can change the event query time in Event Viewer from the last 10 minutes to another value by using the dropdown list.
IPS Missed Packet Trends	A measure of the IPS missed packets trends. The missed packets trend data will appear in this widget only when there is an alert based on missed packets, and the data will disappear when the alert is cleared. You can set monitoring parameters on the IPS Health Monitor page at [IPS device in Device View] Platform > Device Admin > Health Monitor.
IPS License	Displays IPS devices for which the license will expire in 30 days or 60 days. (Use the dropdown list to choose 30 days or 60 days.) If the license will expire in 30 days or 60 days (whichever you select) this widget displays the license expiry date.
IPS Update Packages	Displays sensor updates and signature updates that are available on Cisco.com or on a local download server but not downloaded to the Security Manager server. If there are many such updates, then this widget displays only the 10 most recent updates.
IPS Sensors Out of Date	Sensors requiring a signature update.
Top 10 Reports for Firewall Sources, Destinations, and Services	Pre-defined system reports that you can use to identify the top destinations, services, and sources for firewall ACL events. The statistics are based on the events collected by the Event Manager service (as displayed in Event Viewer). To use these reports, use Report Manager (Launch > Report Manager...).

Table 72-2 Description of Individual WidgetsDashboard widgetswidgets for IPS in Dashboardwidgets for firewall in dashboard (continued)

Top 10 Reports for Botnet Malware Sites, Ports, and Hosts	Pre-defined system reports that you can use to analyze botnet traffic filtering. The statistics are based on the botnet events collected by the Event Manager service (as displayed in Event Viewer) for blocked and gray-listed sites. To use these reports, use Report Manager (Launch > Report Manager...).
Firewall CPU Usage Trends	A measure of the firewall CPU usage trends. The CPU usage trend data will appear in this widget only when a firewall issues an alert because of CPU usage, and the data will disappear when the alert is cleared.

Table 72-2 Description of Individual WidgetsDashboard widgetswidgets for IPS in Dashboardwidgets for firewall in dashboard (continued)

Device Health Summary	Shows current high- or medium-severity active alerts generated by HPM. Alerts can be grouped by Alert-Description, Predefined-Category, Device, or Alert Technology. **Clickable Link**—The device name is an active hyperlink; click it to display the Device Summary dialog box in the Dashboard. This link works for any option in the **Group by** _____ dropdown list: Alert, Category, Device, or Technology. To configure these alerts, use HPM (Launch > Health and Performance Monitor...). **Note** After enabling a device for monitoring in HPM, it can take up to 5 minutes for priority devices and 10 minutes for non-priority devices before actual values can be seen in the Device Health Summary. **Acknowledge Alert**—To acknowledge an alert, follow these steps: 1. Use the **Group by** _____ dropdown list to choose Alert, Category, Device, or Technology. 2. Expand the alert, category, device, or technology that you are interested in. Doing this will show you an alert, category, device, or technology for each device that you are monitoring in Security Manager. 3. Click the **Details** icon (pictured at the end of this topic). Doing this will open the Alert dialog box. 4. Click **Acknowledge Alert** **Clear Alert**—To clear an alert, follow these steps: 1. Use the **Group by** _____ dropdown list to choose Alert, Category, Device, or Technology. 2. Expand the alert, category, device, or technology that you are interested in. Doing this will show you an alert, category, device, or technology for each device that you are monitoring in Security Manager. 3. Click the **Details** icon (pictured at the end of this topic). Doing this will open the Alert dialog box. 4. Click **Clear Alert** You can also access device health summary information from mobile devices. To do this, use the CSM Mobile application. The information available to you from CSM Mobile is the same as that available in the Device Health Summary widget in the Dashboard. For information on enabling or disabling CSM Mobile, see CSM Mobile Page, page 11-9.

Table 72-2 Description of Individual WidgetsDashboard widgetswidgets for IPS in Dashboardwidgets for firewall in dashboard (continued)

Memory Usage Trends	A measure of IPS health status or firewall health trends. For IPS devices, you can set monitoring parameters on the IPS Health Monitor page at [IPS device in Device View] Platform > Device Admin > Health Monitor.
Deployment	Shows the deployment status for all devices for the past 24 hours You can also monitor deployment status by using Deployment Manager (Configuration Manager > Manage > Deployments...) .
IP Intelligence	Information about an IP address related to the following things: • IP Geolocation • FQDN through DNS reverse lookup • WHOIS information For IP Intelligence settings in Security Manager, navigate to Configuration Manager > Tools > Security Manager Administration > IP Intelligence Settings.
CSM Monitor	Shows server information in three categories: • CSM Server Statistics. This information is self-explanatory; for example, the operating system boot time is listed. • CSM User Related Information. This information consists of only one item, the number of users logged in. • CSM DB Backup Related Information. This information tells you if the CSM Monitor widget has found a dangling backup lock file. Knowing if you have a dangling backup lock file is important for the following reason: When a CSM backup is performed, it fails with an error similar to this: "Backup failed.ERROR(383): C:\PROGRA~2\CSCOpx\backup.LOCK file exists." The solution can be described as follows: Security Manager creates a new lock file (backup.LOCK) in the backup directory before it starts a backup. If a backup is interrupted or fails, the file does not get cleaned up. You must delete the current backup.LOCK file from the Security Manager server, and then execute the backup process again. The CSM Monitor widget makes it faster and more convenient for you to detect a dangling backup lock file. For detailed information, refer to the Cisco TAC document at the following URL: http://www.cisco.com/en/US/products/ps6498/products_tech_note09186a0080c13cdd.shtml

Note In some cases, Top Infected Hosts, for example, the dashboard report has a slightly different appearance than the report generated by Report Manager. This is caused by a difference in sorting, but the data is identical. Such a case will occur when more than one entry in the dashboard report has the same count.

Basic dashboard operations are listed in the following table:

Table 72-3 Basic Dashboard OperationsDashboard operations

Launch Dashboard	Configuration Manager or other Security Manager client application > Launch > Dashboard...
Add a new dashboard	File > New Dashboard
Re-arrange Dashboard Tabs for a Default View	You can re-arrange the dashboard tabs so that you can set a default view. For example, you might want the IPS tab to be first (on the extreme left): 1. Click a tab that you are interested in, e.g., Summary, Firewall, or IPS. 2. While your tab of interest is still selected, right-click to see the following context menu options: Move to Left, Move to Right, Move to First, Move to Last. 3. Click your choice. 4. You do not need to save your changes, and your changes will be persistent--the individual dashboard tabs will be arranged in the same way the next time you launch the Dashboard.
Display a different dashboard	Click the tab for the desired dashboard, such as Summary, Firewall, or IPS.
Show or hide widgets	File > Show Widgets or File > Hide Widgets
Add a widget	Drag-and-Drop Method: 1. File > Show Widgets 2. Drag and drop the desired widget onto the dashboard. Menu Method: 1. File > Show Widgets 2. Select the desired widget by clicking it. 3. Click **Add** in the Description bar. 4. Click **Done** in the Description bar. **Note** When using the menu method, the widget will be added to the upper left-hand corner of the Dashboard. If desired, you can rearrange all the widgets by dragging and dropping them.
Remove a widget	Click the **Remove** icon in the title bar of the widget that you want to remove.

Table 72-3 Basic Dashboard OperationsDashboard operations (continued)

Expand a widget	If a widget is shown in the dashboard, you can expand it with the down arrow. The down arrow becomes visible when you hover the mouse pointer over the right side of the widget title bar. The tooltip for the down arrow is labeled "Expand." **Note** A special consideration applies when you 1) collapse a widget, 2) exit the dashboard, and then 3) launch the dashboard again. In this case, you will find that the widget is still collapsed, but the down arrow (normally used to expand it) is not present; only the up arrow (normally used to collapse it) is present. To expand the widget in this case, click the up arrow, note that the down arrow re-appears, and then click the down arrow as usual.
Collapse a widget	If a widget is shown in the dashboard, you can collapse it with the up arrow. The up arrow becomes visible when you hover the mouse pointer over the right side of the widget title bar. The tooltip for the up arrow is labeled "Expand" (not "Collapse").
Group by _____ (Device Health Summary widget only)	A dropdown list offering you the following choices: • Group by Alert • Group by Category • Group by Device • Group by Technology **Note** In the Group by _____ dropdown list, you can click the display name of a device (underlined to indicate that it is a hyperlink) to see an information box on the health of that device in terms of memory and other parameters. The information box contains an address field; the address can be either Host.Domain or the IP address: if Host.Domain is configured, that information will be displayed; otherwise, the IP address will be displayed.

In the Dashboard, many of the icons can be clicked to accomplish a particular action, such as "Refresh" or "Add Dashboard." Most of these "clickable" icons have a tooltip to document the action that clicking the icon will accomplish, but a few do not. Clickable icons that have no tooltip in the dashboard are documented in the following table.

Table 72-4 Clickable Icons that have no Tooltip in the DashboardDashboard iconsicons without tooltips in Dashboard

Icon	Appearance	Widget	Description
	Black exclamation point (bang) on a yellow background in a triangle.	Deployment widget	**Deploying** icon. Indicates that a job is in the deploying state. Click this icon to open/close the job description: • Created date and time • Job Name • Description • State • User • Job Type
	White rectangle (a document) with red and yellow dots.	Deployment widget	**Status Report** icon. Click this icon to see a detailed Deployment Status report
	White checkmark in a green circle with a grey border.	Deployment widget	**Succeeded** icon. Indicates that a job is in the success state. Click this icon to open/close the job description: • Created date and time • Job Name • Description • State • User • Job Type
	White "X" in a red circle with a grey border.	Deployment widget	**Failed** icon. Indicates that a job is in the failed state. Click this icon to open/close the job description: • Created date and time • Job Name • Description • State • User • Job Type

Table 72-4 Clickable Icons that have no Tooltip in the DashboardDashboard iconsicons without tooltips in Dashboard (continued)

Icon	Appearance	Widget	Description
	Clipboard with a pencil making annotations on a sheet of paper.	Device Health Summary widget	**Details** icon. Click this icon to open/close the job description: • Created date and time • Job Name • Description • State • User • Job Type

CSM Mobile

Beginning with Version 4.5, Cisco Security Manager has an application called CSM Mobile.

CSM Mobile allows you to access device health summary information from mobile devices. The information available to you in this way is the same as that available in the Device Health Summary widget in the Dashboard: current high or medium severity active alerts generated by Health and Performance Monitor. Alerts can be grouped by Alert-Description, Predefined-Category, Device, or Alert Technology. For more details on device health summary information in the Dashboard, see Dashboard Overview, page 72-1.

The principal users of CSM Mobile are expected to be those who use an Apple iPad, an Apple iPhone, the Google Chrome browser, or the Apple Safari browser.

CSM Mobile must be enabled for you to use it. For information on enabling or disabling CSM Mobile, see CSM Mobile Page, page 11-9.

Note If the CSM Mobile feature is not enabled, you will be redirected to the default Security Manager login page (which is provided by the CiscoWorks Common Services framework software); you will not receive an error message.

The home page for CSM Mobile has the following alert categories:

- Device Not Reachable
- Interface Down
- Overall Device Health Alerts
- High Memory Utilization
- Firewall—High CPU Utilization
- IPS—High Inspection Load
- IPS—High Missed Packets
- IPS—Bypass Mode
- Other Alerts

Navigation and other tasks in CSM Mobile are accomplished by using a few simple screens and icons:

- **Login**—A screen reading "Cisco Security Manager Mobile—Version 4.5.0" with fields for username and password and a button for login.
- **Logout**—On the CSM Mobile home page, a white X icon on a blue background. This icon is located in the upper left-hand corner.
- **Refresh**—On the CSM Mobile home page, a white circular arrow icon on a blue background. This icon is located in the upper right-hand corner.
- **Alert Detail**—For each type of alert on the CSM Mobile Home page, an grey arrow icon to the right of the alert count.
- **CSM Mobile Back Button**—On each alert detail page, a white angle arrow on a blue pentagon-shaped background [available only on alert detail pages]. The CSM Mobile back button is functionally equivalent to your browser's back button.

Note The CSM Mobile display does not refresh itself automatically; you must manually click the refresh button to obtain up-to-date alert data.

Viewing Inventory Status

You can view a summary of device properties for all devices that you are authorized to view. The summary includes device contact information and all device configurations, indicating which settings are local and which are using a shared policy, and indicating any policy object overrides in effect. You can also view the status of configuration deployment to the device.

The report is in table format, allowing you to organize information by filtering, sorting, reordering and removing columns. You can also export the table contents to a comma-separated values (CSV) file on the Security Manager server.

Step 1 In Device view, select **Tools > Inventory Status** to open the Inventory Status Window, page 72-13.

Step 2 Select the device whose detailed status you want to view in the upper table. The detailed information is shown in the tabs in the lower pane. The information is organized into folders; click the +/- icons to open and close folders, or double-click the folder name. The following tabs are available:

- Inventory—Lists summary information about the selected device's device properties, deployment methods, device group membership, and the parent device for modules.
- Policy—Lists the current status of the policies that can be configured for the selected device, whether the policy is unassigned (not defined), a local policy, or a shared policy.
- Policy Object Overrides—Lists policy objects that have overrides defined for the selected device.
- Status—Lists status messages from Security Manager deployment jobs for the selected device, organized by event type.

  An **event** is a notification that a managed device or component has experienced an abnormal condition. Multiple events can occur simultaneously on a single monitored device or service module.

  Security Manager displays only the most-recent event of each type. To view historical status information, use the Deployment Manager.

Step 3 Click **Close** to close the Inventory Status window.

Inventory Status Window

Use the Inventory Status window to view device properties and status for the devices that you are allowed to view. This window summarizes device information so that you do not have to open the device properties for each individual device.

In addition to device property information, you can view summary information about how the policies on each device are configured (whether local, shared, or not configured) and the policy objects that have overrides for each device. You can also view the status of configuration deployment to the device.

The Inventory Status window contains two panes. Use the upper pane to view a complete listing of all devices, to sort the devices by attribute, or to filter out certain ones. Use the lower pane to view the device property details of the device selected in the upper pane.

Navigation Path

Select **Tools > Inventory Status**.

Related Topics

- Viewing Inventory Status, page 72-12
- Filtering Tables, page 1-48
- Table Columns and Column Heading Features, page 1-49

Field Reference

Table 72-5 Inventory Status Window

Element	Description
Device Summary Information for All Devices (Upper Pane)	
Export button	Click this button to export the inventory as a comma-separated values (CSV) file. You are prompted to specify a file name and to select a folder on the Security Manager server. You can use the export file for reference or analysis.
Display Name	The name of the device as it is displayed in Security Manager.
Deployment	The status of the configuration deployment for the device.
OS Type	The family of the operating system running on the device, for example, IOS, IPS, ASA, FWSM, or PIX.
Running OS Version	The version of the operating system running on the device.
Target OS Version	The target OS version for which you want to apply the configuration. Configurations are based on the commands supported by this version.
Host Name.Domain Name	The DNS host and domain names for the device.
IP Address	The management IP address of the device.
Device Type	The type of device.
Details for the Selected Device (Lower Pane)	
The detailed information is shown in the tabs in the lower pane. The information is organized into folders; click the +/- icons to open and close folders, or double-click the folder name.	

Table 72-5 Inventory Status Window (continued)

Element	Description
Inventory	Lists summary information about the selected device's device properties, deployment methods, device group membership, and the parent device for modules.
Policy	Lists the current status of the policies that can be configured for the selected device, whether the policy is unassigned (not defined), a local policy, or a shared policy.
Policy Object Overrides	Lists policy objects that have overrides defined for the selected device. For more information on policy object overrides, see Policy Object Override Pages, page 3-52.
Status	Lists any deployment status messages for the selected device. Events are organized by event type. Event details include timestamp, description, and recommended action. The time stamp indicates the time of the last change in status for the device, not the time of the latest polling of the device. Also shown is the highest severity level of the status messages.
Navigation buttons	Click the navigation buttons to move through the inventory list. From left to right, buttons mean go to the first device in the list, go to the previous device, go to the next device, and go to the last device. The center field indicates which device is currently selected based on the row number (for example 5/10 means the fifth of 10 devices in the list).

Starting Device Managers

You can start a device manager to view a device's configuration and status from within Security Manager. You can start device managers for ASA, ASA-SM, PIX, FWSM, IPS, and IOS devices.

Each device manager includes several monitoring and diagnostic features that provide information regarding the services running on the device and a snapshot of the overall health of the system. You can use these device managers to view the existing device configuration and to monitor current status, but you cannot use it to apply configuration changes to the device.

Note Following points are to be noted:

- You cannot start device managers for IPS virtual sensors.
- In Cisco Security Manager 4.16, due to upgrade of JRE 1.7 build 161, support to some of old applets are dropped. Hence, beginning from Cisco Security Manager 4.16 you cannot launch PIX 6.3, IDS/IPS versions 5.x to 7.x, and FWSM 2.x directly.
- To cross-launch ASDM from Cisco Security Manager 4.18 for ASA 9.10.1(10) and later versions, you must enable the "`http server basic-auth-client Java`" CLI manually in the device.

To start a device manager, select the device in Device view, right-click and select **Device Manager**. You can also start the device manager by selecting **Launch > Device Manager**. (These commands are disabled when you select an ASA CX device, and the **Prime Security Manager** commands are enabled. Cisco Prime Security Manager is used to configure and manage ASA CX devices. See Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20 for more information.)

When you start a device manager from Security Manager, the device manager executable is downloaded to your client system; the device manager does not need to be installed on the network device. The first time you start a device manager, it takes time to download the software to your workstation (you are shown a progress bar). (If you run into problems, review the tips in Troubleshooting Device Managers, page 72-16.)

Security Manager selects the most appropriate device manager version based on the operating system running on the network device. Subsequent communications with the selected device are completely transparent. Connections are made through the Security Manager server; that is, the Security Manager server acts as a proxy server. By starting a device manager from Security Manager, you eliminate the need to open an HTTPS connection between your client system and the device you want to monitor.

Tip When you start a device manager session, Security Manager opens a version of the manager that is appropriate for the operating system software version running on the device. However, Security Manager might not open the most recently-available version of the device manager if new device manager versions have been released after the release of the Security Manager version you are using. When you start the device manager, check its version (for example, select **Help > About** in the device manager window); if there is a more recent device manager available with features that you require, you must install and use that device manager outside of Security Manager to use those new features.

Keep in mind that if you use an external device manager running on the device to modify device configurations directly, these changes are considered out-of-band by Security Manager, and might be subsequently overwritten when you next deploy configurations from Security Manager. For more information about out-of-band changes, and what you can do to identify and recreate them, see the following topics:

- Understanding How Out-of-Band Changes are Handled, page 8-12
- Detecting and Analyzing Out of Band Changes, page 8-45

Security Manager starts only one instance of a device manager per device, and closes the device manager when you exit Security Manager, or when the idle-session timeout period is exceeded. You can have more than one device manager window open at one time (connected to different devices).

The following table outlines the device managers you can launch from Security Manager.

Table 72-6 Device Managers Available in Security Manager

Device Manager	Description
IDM	The IPS Device Manager (IDM) lets you monitor IPS sensors and modules that are part of the Security Manager inventory. See the IDM documentation for more information about using this device manager.
PDM	The PIX Device Manager (PDM) lets you monitor PIX 6.x devices and early FWSMs, specifically FWSM releases 1.1, 2.2 and 2.3 in single- or multiple-context modes. See the PDM documentation for more information about using this device manager.

Table 72-6 Device Managers Available in Security Manager (continued)

Device Manager	Description
ASDM	The Adaptive Security Device Manager (ASDM) lets you monitor ASA, ASA-SM, PIX 7.x+, and FWSM 3.x+ devices. See the ASDM documentation for more information about using this device manager.
SDM	The Security Device Manager (SDM) lets you monitor Cisco IOS-based resources. SDM requires no previous experience with Cisco devices or the Cisco command-line interface (CLI). Cisco SDM supports a wide range of Cisco IOS software releases. See the SDM documentation for more information about using this device manager.

The following topics explain more about troubleshooting and using device managers:

- Troubleshooting Device Managers, page 72-16
- Access Rule Look-up from Device Managers, page 72-17
- Navigating to an Access Rule from ASDM, page 72-18
- Navigating to an Access Rule from SDM, page 72-19

Troubleshooting Device Managers

If you can successfully deploy configurations to a device, Security Manager should be able to open a device manager session with the device (as described in Starting Device Managers, page 72-14). For cross-launching the Device Manager for ASA 9.10.1(10) and later versions, you must enable the "`http server basic-auth-client Java`" CLI manually in the device. However, if you have problems making a connection or using one that is open, consider the following troubleshooting tips, which are divided into basic tips and tips for using multiple device managers.

Basic Device Manager Troubleshooting Tips

- Generally, the credentials configured for the device in the Security Manager inventory are used to start the device manager. However, some versions of SDM require that you enter a user name and password when the device manager is started. If you get an error that says device credentials are missing, or they are not valid, update the Device Properties Credentials page with a username and password that can log into the device. In Device view, right-click the device and select **Device Properties**. For more information, see Viewing or Changing Device Properties, page 3-41 and Device Credentials Page, page 3-45.
- All users associated with any of the CiscoWorks Common Services roles have permission to start device managers from Security Manager, with the exception of the Help Desk role or any of the predefined Cisco Secure ACS roles. Ensure you have appropriate permissions.
- SSL/HTTPS must be enabled on the target device to provide secure communications between Security Manager and the device. An error message is displayed if SSL is not enabled on the device. See Understanding Device Communication Requirements, page 2-1 for more information.
- You might need to modify Cisco Security Agent, or other anti-virus and network firewall software, on the Security Manager system and on your workstation to allow the device manager service (**xdm-launcher.exe**) to be started.

- Ensure that Security Manager is correctly configured for contacting and communicating with the target device. Specifically verify device properties such as identity, operating system and credentials. Select the desired device, right-click and choose **Device Properties**. Verify the settings on the General and Credentials pages. You can test whether Security Manager can connect to the device by selecting the Credentials tab and clicking **Test Connectivity** (see Testing Device Connectivity, page 9-1).
- Device managers can be started for FWSMs and ASAs running in transparent mode (Layer 2 firewall) or routed mode (Layer 3 firewall), and supporting a single security context or multiple security contexts. For FWSM and ASA devices running multiple security contexts, you must define a unique management IP address for each security context.
- If you get a message saying that the platform is not supported for device manager launch, but you believe the platform should be supported based on information in this guide, consider the relative newness of the operating system version running on the device and the age of the Security Manager software version you are using. If you are using very recent operating systems, but a relatively older version of Security Manager, you might need to upgrade Security Manager (or apply a service pack), contact Cisco Technical Support, or simply install the latest device manager on the network device and use it outside of Security Manager. Before using a device manager outside of Security Manager, review the information on out-of-band changes in Starting Device Managers, page 72-14.

Multiple Device Manager Sessions Troubleshooting Tips

- Starting multiple device managers might affect the performance of both the Security Manager server and your client. On the client, memory requirements and performance impact are proportional to the number of device managers launched. On the server, a large number of requests to start device managers or retrieve current information from the device can have an adverse impact on performance.
- The maximum number of persistent HTTPS connections that can be established with any one device from all clients depends on the device type and model. An error message is displayed if you attempt to exceed this limit.

  For example, a single PIX 6.x allows multiple clients to each have one browser session open, supporting up to 16 concurrent PDM sessions. An FWSM (1.1, 2.2, or 2.3) allows up to 32 PDM sessions for the entire module, with a maximum of five concurrent HTTPS connections per context.

  Refer to the appropriate device documentation for information about specific limits.

Access Rule Look-up from Device Managers

A set of access rules is associated with each device interface. These rules are presented in the form of an ordered list or table. This list is often referred to as an access-control list (ACL), with each rule in the list known as an access-control entry (ACE). When deciding whether to forward or drop a packet, the device tests the packet against each access rule in the order listed. When a rule is matched, the device performs the specified action, either permitting the packet into the device for further processing, or denying entry. If the packet does not match any rule, the packet is denied.

Activity on your firewall or router can be monitored through syslog messages. If logging is enabled on the device, whenever an access rule that is configured to generate syslog messages is matched—for example, a connection was attempted from a denied IP address—a log entry is generated.

Note For the device to generate log entries, logging must be enabled on the device (on the Logging Setup Page for ASA/PIX devices and the Logging policies for IOS devices, described in Logging on Cisco IOS Routers, page 65-1), and the individual access rules must be configured to generate log messages when they are matched (in the Advanced and Edit Options Dialog Boxes, page 16-17).

You can monitor syslog messages in device managers launched from Security Manager. For certain device managers, you can also look up the access rule in Security Manager that generated a particular message from the monitoring window. The access rule that triggered the syslog entry is highlighted in Security Manager on a first-match basis, even if there are multiple matches.

This access rule look-up is available through SDM for all managed routers running IOS, and through ASDM for managed PIX and ASA devices (including ASA-SM) running version 8.0(3) and above, and FWSM devices running version 3.1 and above.

The following topics describe how to look up access rules in Security Manager from a device manager:

- Navigating to an Access Rule from ASDM, page 72-18
- Navigating to an Access Rule from SDM, page 72-19

Navigating to an Access Rule from ASDM

In an ASDM device manager launched from Security Manager, you can monitor system log messages in the Real-time Log Viewer window and the Log Buffer window. You can select a syslog message displayed in either window and navigate to the access-control rule in Security Manager that triggered the message, where you can update the rule as necessary.

The Real-time Log Viewer is a separate window that lets you view syslog messages as they are logged. The separate Log Buffer window lets you view messages present in the syslog buffer.

You can look up access rules associated with the following syslog message IDs:

- 106023 – Generated when an IP packet is denied by the access rule. This message appears even when logging is not enabled for the rule.
- 106100 – If logging is enabled for a matched access rule (in the Advanced and Edit Options Dialog Boxes, page 16-17), this message provides information about the traffic flow, depending on the parameters set. This message provides more information than message 106023, which logs only denied packets.

This procedure describes how to look up an access rule in Security Manager from ASDM's Real-time Log Viewer or Log Buffer windows.

Related Topics

- Access Rule Look-up from Device Managers, page 72-17
- Navigating to an Access Rule from SDM, page 72-19

Step 1 Select a PIX, ASA, ASA-SM, or FWSM in the Security Manager device inventory.

Step 2 Select **Launch > Device Manager** to start ASDM. For more information about starting device managers, see Starting Device Managers, page 72-14.

Step 3 In the ASDM window, click the **Monitoring** button to display the Monitoring panel; click **Logging** in the left pane to access the log-viewing options.

Step 4 Select either **Real-time Log Viewer** or **Log Buffer**.

Step 5 Click the **View** button to open the selected log-viewing window.

Note The View button is not displayed if logging is not enabled on the device.

Each syslog message listed in the window includes the following information: message ID number, date and time the message was generated, the logging level, and the network or host addresses from which the packet was sent and received.

Step 6 To view the access rule that triggered a specific syslog message, select the message and click the **Show Rule** button in the ASDM toolbar (or right-click the message and choose **Go to Rule in CSM** from the pop-up menu).

The Security Manager client window is activated and the Access Rules page appears with the rule highlighted in the rules table. If the syslog entry was triggered by an access rule not referenced in the current Security Manager activity, an error message appears.

Navigating to an Access Rule from SDM

In an SDM device manager launched from Security Manager, you can view a log of events categorized by security level under the Syslog tab of the Logging window. You can select a syslog message and navigate to the access-control rule in Security Manager that triggered the message, where you can update the rule as necessary.

The Monitor > Logging option in SDM offers four log tabs; Syslog is the only one of these offering the Security Manager access-rule look-up option. The router contains a log of events categorized by severity level. The Syslog tab displays the router log, even if log messages are being forwarded to a syslog server.

On Cisco IOS devices, syslog messages are generated for access rules configured with the **log** or **log-input** keywords. The **log** keyword produces a message when a packet matches the rule. The **log-input** keyword produces a message that includes ingress interface and source MAC address, in addition to the packet's source and destination IP addresses and ports. When identical packets are matched, the message is updated at five-minute intervals with the number of packets permitted or denied in the previous five minutes.

This procedure describes how to look up an access rule in Security Manager from the Syslog tab of SDM's Logging panel.

Related Topics

- Access Rule Look-up from Device Managers, page 72-17
- Navigating to an Access Rule from ASDM, page 72-18

Step 1 Select an IOS router in the Security Manager device inventory.

Step 2 Select **Launch > Device Manager** to start SDM. For more information about starting device managers, see Starting Device Managers, page 72-14.

Step 3 In the SDM window, click the **Monitoring** button to display the Monitoring panel; click **Logging** in the left pane to access the log-viewing options.

The Logging pane appears with Syslog tab displayed.

Step 4 To view the access rule that triggered a specific syslog message, select the message and click the **Go to Rule in CSM** button above the table of log messages.

The Security Manager client window is activated and the Access Rules page appears with the rule highlighted in the rules table. If the syslog entry was triggered by an access rule not referenced in the current Security Manager activity, an error message appears.

Launching Cisco Prime Security Manager or FireSIGHT Management Center

The ASA CX is an Adaptive Security Appliance module that provides advanced ConteXt-aware security, extending the ASA platform to provide in-depth "who-what-where-when-how" application visibility and control. The ASA FirePOWER module supplies next-generation firewall services, including Next-Generation IPS (NGIPS), Application Visibility and Control (AVC), URL filtering, and Advanced Malware Protection (AMP).

The ASA CX devices are managed by the Cisco Prime Security Manager (PRSM) application and the ASA FirePOWER modules are managed by the FireSIGHT Management Center application—they cannot be directly managed by Cisco Security Manager. However, Security Manager has been enhanced to allow you to discover the presence of these modules on ASA devices; to "cross launch" PRSM and FireSIGHT Management Center from the Configuration Manager application; and to share Policy Object data between Security Manager and PRSM.

Note PRSM and FireSIGHT Management Center are browser-based applications; that is, they are launched and operate within a browser window. Therefore, when you cross-launch PRSM or FireSIGHT Management Center from the Configuration Manager client, the host system's default browser is opened and the management application initiated. However, some browsers have not been certified with PRSM or FireSIGHT Management Center and you may need to change the default browser on the Security Manager client's host system prior to cross-launching. See "Browser Requirements" in the PRSM or FireSIGHT Management Center installation guide for more information.

Before You Begin

In order to cross-launch PRSM or FireSIGHT Management Center, Security Manager must be aware of the presence of the modules. This is accomplished through discovery of either new ASA devices, or of modules which have been added to existing ASAs. This process is outlined in Detecting ASA CX and FirePOWER Modules, page 72-21.

Also, you can enable and configure "single sign-on" (SSO) to allow Security Manager users direct access to PRSM or FireSIGHT Management Center without logging into the applications separately. To allow this, the appropriate user credentials must be defined in both applications. (Note that SSO is not necessary to cross-launch PRSM or FireSIGHT Management Center.) See Security Manager's Single Sign-on Configuration Page, page 11-68, and "Configuring Single Sign-On for Cisco Security Manager" in the *User Guide for ASA CX and Cisco Prime Security Manager* (Cisco ASA CX Context-Aware Security End-User Guides) for more information.

Related Topics

- Single Sign-on Configuration Page, page 11-68
- Detecting ASA CX and FirePOWER Modules, page 72-21
- Sharing Device Inventory and Policy Objects with PRSM, page 72-22

To monitor and manage your ASA CX devices or FirePOWER Modules, cross-launch PRSM or FireSIGHT Management Center:

Step 1 Select a previously discovered ASA CX device or ASA with FirePOWER module in Configuration Manager's Device view—in either the device-selector tree, or the table of devices in the content area.

Again, discovering ASA CX devices or FirePOWER modules in Security Manager is described in Detecting ASA CX and FirePOWER Modules, page 72-21.

Step 2 Right-click the selected device and choose **Prime Security Manager** or **FireSIGHT Management Center** from the pop-up menu. Alternatively, you can choose **Prime Security Manager** or **FireSIGHT Management Center** from the Configuration Manager's **Launch** menu. (These commands are available only when you have selected an ASA CX or an ASA with a FirePOWER module.)

The browser-based PRSM or FireSIGHT Management Center window appears, displaying the device screen for the selected device.

Note The URL used by Security Manager to launch PRSM incorporates the management IP address of the CX module (obtained during device detection), and includes the string `/admin/mgmt?rtp`. During cross-launch, this type of request is redirected to the appropriate PRSM central server, if one exists. Otherwise, the "on-box" version of PRSM is launched. (To directly launch the on-box version of PRSM yourself, you must type ***https://<management_IP_address>***, where *<management_IP_address>* is the management address of the desired CX module, into your browser's address field.)

Information about using PRSM can be found on the Cisco ASA CX Context-Aware Security End-User Guides page of cisco.com and information about using FireSIGHT Management Center can be found on the Cisco FireSIGHT Management Center page of cisco.com.

Detecting ASA CX and FirePOWER Modules

Prior to being able to share Policy Object data between Security Manager and PRSM, and to cross-launch PRSM or FireSIGHT Management Center from Configuration Manager, you must ensure that Security Manager is aware of the module.

Detection of a CX or FirePOWER module is automatic when you discover a new ASA device by selecting the relevant options in the New Device wizard, as described in Adding Devices to the Device Inventory, page 3-7.

When you add a CX or FirePOWER module to an ASA device already in the inventory, you can detect the new module without affecting the existing policies on the host ASA, as follows:

1. Select one or more ASA devices in Configuration Manager's device-selector tree.

   You can detect more than one module at once—any selected devices that are not ASAs, or that are ASAs which do not include a CX or FirePOWER module, are ignored.

2. Right-click any selected device and choose **Detect ASA-CX/FirePOWER Module** from the pop-up menu.

   The Create Discovery Task dialog box or Bulk Rediscovery dialog box appears, with the *Detect ASA-CX/FirePOWER Module* option selected—none of the other discovery options are available.

See Create Discovery Task and Bulk Rediscovery Dialog Boxes, page 5-18 for more information about using this dialog box.

3. Click **OK** on the Create Discovery Task dialog box or click **Finish** on the Bulk Rediscovery dialog box to close the dialog box and begin module detection.

   You may be warned that discovery will replace existing policies; you can safely click Yes to close the warning and proceed.

   The Discovery Status dialog box opens automatically displaying discovery progress; see Viewing Policy Discovery Task Status, page 5-22 for more information about this process.

When a CX or FirePOWER module is detected on an ASA, the management IP address of the module itself is fetched and the ASA-CX/FirePOWER Module section of the Device Properties window is updated; see Device Properties: General Page, page 3-41. The management IP address is used to cross-launch PRSM or FireSIGHT Management Center. (Cisco Prime Security Manager, or PRSM, is the application used to configure and manage ASA CX devices and FireSIGHT Management Center is the application used to configure and manage ASA FirePOWER modules, as described in Launching Cisco Prime Security Manager or FireSIGHT Management Center, page 72-20.)

Note The URL used by Security Manager to launch PRSM incorporates the management IP address of the CX module (obtained during device detection), and includes the string `/admin/mgmt?rtp`. During cross-launch, this type of request is redirected to the appropriate PRSM central server, if one exists. Otherwise, the "on-box" version of PRSM is launched. (To directly launch the on-box version of PRSM yourself, you must type ***https://<management_IP_address>***, where *<management_IP_address>* is the management address of the desired CX module, into your browser's address field.)

Upon completion of the detection process, all ASAs with CX modules installed are indicated in the various Security Manager displays by presentation or inclusion of the PRSM icon: . For example, here is the ASA CX icon used in the Device selector: .

Caution You also can detect the presence of a CX or FirePOWER module on an existing ASA by choosing **Discover Policies on Device(s)** from the selected-device right-click menu, or by choosing **Discover Policies on Device** from the Policy menu. Depending on the number of devices selected and which command you choose, the Create Discovery Task dialog box, or the Bulk Rediscovery Task dialog box, opens and all discovery-rediscovery options are available. This means you can potentially overwrite any shared policies already established on the selected device(s). Be sure to deselect all options except **Detect ASA-CX/FirePOWER Module**, unless you are sure you want to rediscover existing policies. See Discovering Policies on Devices Already in Security Manager, page 5-15 for more information.

Sharing Device Inventory and Policy Objects with PRSM

You can export the current device inventory, and the current set of policy objects, as defined in Security Manager for import into Cisco Prime Security Manager (PRSM).

Exporting the Device Inventory

To share the Security Manager device inventory with PRSM, export the inventory as a comma-separated values (CSV) file, as described in Exporting the Device Inventory, page 10-6. Be sure to specify "Cisco Security Manager" as the format type for the export file.

Exporting Networks/Hosts and Services Policy Objects

To export Security Manager policy objects—specifically Networks/Hosts objects, or Services objects; PRSM does not support Port List objects—for import into PRSM, you must execute a Perl script on the Security Manager server host to create a CSV file.

The Perl script is included in the Security Manager server installation, and its use is described in detail in Importing and Exporting Policy Objects, page 6-23. The basic procedure is as follows:

1. Log into the computer running the Security Manager server, open a Cmd window, navigate to the Perl-script location, and then execute the Perl-script command at the command prompt.

   Here is an example of the command as used to export Networks/Hosts objects: `perl PolicyObjectImportExport.pl -u` *`user`* -p *password* -o export -t network -f C:\CSM_Net_objects.csv -e true

2. Copy the CSV file to the PRSM client system.

   This file can be edited, if necessary.

3. Launch PRSM and import the CSV file. For information about this process, see the section "Importing Objects" in the "Managing Policy Objects" chapter of the PRSM user guide.

Analyzing an ASA or PIX Configuration Using Packet Tracer

Note From version 4.17, though Cisco Security Manager continues to support PIX features/functionality, it does not support any bug fixes or enhancements.

Packet tracer is a policy debugging tool for ASA and PIX security appliances running version 7.2.1+ that are operating in router mode.

The packet tracer inspects the active policies currently running on the appliance. Without having to generate real traffic, you can analyze how traffic between two addresses traverses the security appliance, whether it is dropped or allowed. If the result is unexpected, you can determine where the issue exists and update the corresponding policy in Security Manager to resolve it.

Packet tracer presents a step-by-step analysis of how a simulated packet is processed by the security appliance's active configuration. It traces the packet's flow through the active firewall modules, such as route lookup, access lists, NAT translations, and VPN. The set of active modules changes based on the type of packet configured and the active configuration. For example, if no VPN policies are configured, the VPN module is not evaluated.

You can inspect the simulated packet's traversal rather than having to generate network traffic, enable syslog messages, and manually review resulting syslog messages. Packet tracer details the actions enforced by the active configuration on the packet. If a configuration command causes the packet to be dropped, the reason is provided, such as "Drop-reason: (telnet-not-permitted) Telnet not permitted on least secure interface."

You can trace the life span of a simulated packet through the security appliance to see whether the packet is behaving as expected. Packet tracer uses include the following:

- Debug all packet drops in a production network.
- Verify the configuration is working as intended.
- Show all rules applicable to a packet including the CLI that defines the rule.
- Show a time line of packet changes in a data path.

- Trace packets in the data path.
- If the packet is blocked or permitted by some explicit access rule, you can use a short-cut to go to the policy so that you can edit the rule.

Tips:

- Packet Tracer is also available in the ASDM application and the ASA command line, and the Security Manager version is equivalent to the ASDM version. For an example of using Packet Tracer from ASDM and the CLI to analyze a configuration, see PIX/ASA 7.2(1) and later: Intra-Interface Communications.
- Before you can use packet tracer on a device, you must submit your policy changes at least once after adding the device to the inventory.
- Packet tracer analyzes only the active configuration running on a device. Therefore, you cannot use packet tracer to test proposed configurations before they are deployed and running on the device. Do not use packet tracer on a device with pending configuration changes—deploy the changes first and then use packet tracer to ensure the packet tracer results are valid.

To use Packet Tracer:

Step 1 (Device view) Right-click on the ASA or PIX 7.2.1+ device and select **Packet Tracer** on the shortcut menu to open the Packet Tracer window.

Step 2 Select the interface you want to test from the **Interfaces** list. The list contains all interfaces defined on the device.

Step 3 Model the packet that you want to trace by configuring the following fields:

- **Packet Type**—Select whether you are tracing a TCP, UDP, ICMP, IP, or ESP packet.

Note Beginning with 4.16, Cisco Security Manager supports tracing of a ESP packet from ASA 9.9.1 devices.

- **Source, Destination IP Address**—Select from the following address types and enter the host addresses for both ends of the communication (from source to destination):
  - IP Address of the host. This can be IPv4 or IPv6 addresses. Packet tracer with IPv6 is not supported for devices running ASA software version lower than 8.4(2).
  - User (source only). For example, DOMAIN\Administrator. The IP address mapped to the user is used for the trace. You must enable identity-aware firewall by configuring identity options to use this type of address.
  - FQDN, or fully-qualified domain name, of the host. For example, host.example.com. You must configure DNS to use this type of address.
  - Security Name (ASA 9.x+ only).
  - Security Tag (ASA 9.x+ only).
- **Source, Destination Port (TCP and UDP only)**—Enter, or select, the port numbers that represent the traffic type. The selection list uses names that equate to the standard port numbers for the named application. For example, selecting **http** and entering **80** is the same.
- **Type, Code, ID (ICMP only)**—When modeling an ICMP packet, you must enter values in all of these fields:

  - **Type**—Select the ICMP packet type or enter the equivalent number. The list includes all main ICMP types. For a complete list of types and related codes, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for "ICMP Type Numbers."
  - **Code**—Enter **0** unless you are modeling a packet type that has non-zero codes. These are destination unreachable (type 3, codes 0-12), redirect (type 5, codes 0-3), time exceeded (type 11, codes 0-1), and parameter problem (type 12, codes 0-2). See RFC 1700 for code explanations, and note that additional codes might have been introduced in other RFCs.
  - **ID**—You must enter a value for ID even though the field is used for a limited number of message types only. The ID is used for ICMP types that include request and reply versions, such as echo and echo request, to help match replies to requests. The value should be between 1-255.

- **Protocol (IP only)**—Enter the number that identifies the next level protocol. For a complete list of protocol codes, see RFC 1700 at http://www.ietf.org/rfc/rfc1700.txt and search for the "Protocol Numbers" heading. As of the writing of this topic, numbers 1-54 and 61-100 represent values assigned to actual protocols from the accepted range of 0-255.
- **VLAN ID (1- 4096)**—Enter the VLAN ID for the flow. The VLAN ID determines, which VLAN the packet belongs to. Cisco Security Manager validates the ID range to be between 1 and 4096.

Note Beginning with Version 4.13, Cisco Security Manager Packet tracer supports the transparent FW devices. Vlan ID is the new parameters introduced in Version 4.13 to support the packet tracer on devices 9.7.1 onwards.

- **Destination MAC**—Enter the destination MAC address for the flow. Cisco Security Manager validates the MAC address for its format.
- **Enter the SPI (ESP only)**—Enter the security parameter index. It is the arbitrary value used (together with the destination IP address) to identify the security association of the receiving party. Enter a numeric value between 0 and 4294967295.

Step 4 From the **Tracing Packet** drop-down list, select the relevant option:

- bypass-checks—bypasses all security checks for the simulated packet
- decrypted—treats simulated packet as IPsec/SSL VPN decrypted
- persist—enables long term tracing and follow tracing in cluster
- transmit—allows simulated packet to be transmitted from device

Step 5 If you want to see the progress of the trace while it is happening, select **Show animation**. Otherwise, the window is not updated with the results until the trace is completed.

Step 6 Click **Start** to trace the packet.

The policies are examined, and the bottom of the window shows the results in two forms: graphical and detailed information. The graphical view summarizes the phases evaluated in the packet's path. Checkmarks indicate the packet passed the phase, a red X indicates the packet was dropped at that point.

The detailed information organizes the results in folders that correspond to the phases, with an Action column that indicates the results of the phase (checkmark for passed, red X for dropped). To open a folder, click its heading. Detailed information can include the specific configuration commands evaluated and the data derived from **show** commands. The final folder, named Result, summaries the results of the trace.

Tips:

- If the packet is allowed or denied by an explicit access rule, then you can jump to that rule. Select the Access-List folder to open it, then click the **Show access rule** link at the top of the section. You are taken to the Access Rule policy with the rule highlighted; you can edit the rule as desired. If a packet is dropped due to an implicit drop rule, the Show access rule link is not available because the rule does not exist in the policy table.
- If the device is shut down or not reachable due to a network failure during the analysis, an error message stating "Device Connectivity is Failed" appears.
- If you start a new trace, the information shown is cleared automatically. However, you can clear it yourself by clicking **Clear**.

Analyzing Connectivity Issues Using the Ping, Trace Route, or NS Lookup Tools

You can use the Ping or Trace Route tools to investigate and troubleshoot your network configuration and connectivity. You typically run these commands in the device, from within Security Manager by specifying specific launch points and parameters. This causes Security Manager to generate the corresponding command. NS Lookup, on the other hand, is typically run from the Security Manager client.

Note Beginning with Version 4.13, Cisco Security Manager trace routes supports IPv6 address. From ASA version 9.7.1, traceroute for IPv6 address is supported.

Table 72-7 Profiles of the Ping, Trace Route, and NS Lookup Troubleshooting Commands

Tool	Profile
Ping	Use Ping to test whether a particular host is reachable across an IP network and to measure the round-trip time for packets sent from the local host to a destination computer. This can include measuring the local host's own interfaces using ICMP messages. See Analyzing Configuration Using Ping, page 72-27 for details on using this tool.
Trace Route	Use trace route to show the route taken by packets across an IP network. The system returns the number of hops taken and the addresses of each device traversed. See Analyzing Configuration Using TraceRoute, page 72-28 for details on using this tool.
NS Lookup	Use NS lookup (namespace lookup) to issue an NS lookup command from a device so you can test the contents of the DNS server that the queried device uses. See Analyzing Configuration Using NS Lookup, page 72-30 for details on using this tool.

Applicability

The Ping tool is applicable on the following devices: ASA (7.0 – 8.3), PIX [6.3(1-5) to 8.0(2-4)], FWSM [2.2(1) – 4.1(1)], all IOS. It is not applicable to IPS.

The Trace Route tool is applicable on the following devices: ASA [7.2(1) and onward], PIX [6.3(1-5) to 8.0(2-4)], and all IOS. It is not applicable to FWSM nor IPS.

The NS Lookup tool is not supported in any of the devices managed by Cisco Security Manager; rather, you run it from the Cisco Security Manager client using the Windows API.

Analyzing Configuration Using Ping

The ping tool, by default uses the ICMP echo request and echo reply messages to test reachability to a remote system. You can also choose to employ TCP to ping. In its simplest form, ping simply confirms that an IP packet is capable of getting to and getting back from a destination IP address. A ping is sent to an IP address and it returns a reply. This process enables network devices to discover, identify, and test each other. From within Security Manager, you can designate both the network device from which to issue the ping command, and the target of the echo request. This tool generally returns two pieces of information: whether the source can reach the destination (and, by inference, vice versa), and the round-trip time (RTT, typically in milliseconds).

You can use the Ping diagnostic tool in a variety of ways, including:

- **Pinging to a security appliance**—Ping an interface on another security appliance to verify that it is up and responding.
- **Loopback testing of two interfaces**—Initiate a ping from one interface to another on the same security appliance, as an external loopback test to verify basic “up” status and operation of each interface.
- **Pinging through a security appliance**—Ping packets originating from the Ping tool may pass through an intermediate security appliance on their way to a device.The echo packets also pass through two of its interfaces as they return. You can use this to perform a basic test of the interfaces, operation, and response time of the intermediate unit.
- **Pinging to test intermediate communications**—Initiate a ping from a security appliance interface to a network device that is known to be functioning correctly and returning echo requests. If you receive the echo, you confirm physical connectivity and the correct operation of any intermediate devices.

Tip From within the Event Manager, you can right-click on an event to open the Ping Tool and ping the associated device.

Step 1 In Device view, select **Tools > Ping, TraceRoute and NS Lookup . . .**.

The Ping, TraceRoute and NS Lookup dialog box appears.

Step 2 From the device selector, select the device from which to issue the **Ping** command.

The selected device is listed in the top right of the dialog box.

Note To employ TCP for the ping, select TCP for the Packet Type. (The default packet type is ICMP)

Step 3 In Hostname/IPv4address, enter the IP address of the host network/host policy object to be pinged.

Alternatively, click **Select** to choose a host network/host object that defines the host network/host policy object to be pinged.

Step 4 Enter a timeout value. [Optional]

Step 5 Click **Ping**.

The results are displayed in the lower window area.

Example Ping output:

```
Sending 5, 100-byte ICMP Echos to out-pc, timeout is 2 seconds:

!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms
```

Example Ping unsuccessful output:

```
Sending 5, 100-byte ICMP Echos to 10.132.80.101, timeout is 2 seconds:
?????

Success rate is 0 percent (0/5)
```

You can click **Clear Output** to remove the previous response from the Ping output area.

For additional details on the **ping** command, see Troubleshooting TCP/IP on Cisco.com.

Analyzing Configuration Using TraceRoute

The Traceroute tool helps you to determine the route that packets will take to their destination. The tool prints the result of each probe sent. Every line of output corresponds to a TTL value in increasing order.

Traceroute can return useful information about TCP/IP connectivity across your network. The following table shows some of the codes that can be returned by the Traceroute utility, along with their possible cause.

Table 72-8 Traceroute Output Symbols

Output Symbol	Description
*	No response was received for the probe within the timeout period.
nn msec	For each node, the round-trip time (in milliseconds) for the specified number of probes.
!N.	ICMP network unreachable.
!H	ICMP host unreachable.
!P	ICMP unreachable.
!A	ICMP administratively prohibited.
?	Unknown ICMP error.

Step 1 In Device view, select **Tools > Ping, TraceRoute and NS Lookup . . .**.

The Ping, TraceRoute and NS Lookup dialog box appears.

Tip From within the Event Manager, you can right-click on an event to open the TraceRoute page and trace the route of the associated device.

Step 2 Select the **Trace Route** tab.

The Trace Route page appears.

Step 3 From the device selector, select the host from which the route is traced.

Step 4 Enter the **IP Address/Hostname** to specify the address or name of the host to which the route is traced.

Alternatively, click **Select** to choose a host network/host object that defines the IP address.

Note Beginning with Version 4.13, Cisco Security Manager trace routes supports IPv6 address. Till Version 4.12, the Syslog server was configured with Devices having IPv4 address. From Device version 9.7.1, traceroute for IPv6 address is supported. The Syslog server can be configured with IPv6 syslog address with devices having IPv6 addresses.

Step 5 Specify values, as required, for the following:

Table 72-9 Traceroute Fields

Field	Description
Timeout[optional]	The amount of time, in seconds, to wait for a response before the connection times out. The default is three seconds.
Port[optional]	The destination port used by the UDP probe messages. The default is 33434.
Probes per hop[optional]	The number of probes to be sent at each TTL level. The default is three.
TTL Min[optional]	The minimum TTL value for the first probes. (Default is 1.)
TTL Max[optional]	The maximum TTL value for the first probes.(Default is 30.)

Step 6 If desired, select **Specify Source Interface or IP Address**, and then do one of the following:

- Select a source **Interface** from the drop-down list

Note If an IPv6 address is specified in the **IP Address/ Hostname** field, the source **Interface** field is not applicable.

- Enter an **IP Address**

Step 7 If desired, select **Reverse Resolve** to reverse between displaying the address or hostname.

Step 8 If desired, select **ICMP** to use that protocol rather than IP.

Step 9 Click **Trace**.

The traceroute terminates when the packet reaches the destination or when the TTL Max value is reached. The hops taken and the device address corresponding to each hop are displayed.

Analyzing Configuration Using NS Lookup

You use the NS Lookup tool to look up a remote host address when you have the hostname, or to look up the hostname when you have the address.

Unlike the Ping and Traceroute tools, NS Lookup is done on the Security Manager client.

Step 1 In Device view, select **Tools > Ping, TraceRoute and NS Lookup . . .**.

The Ping, TraceRoute and NS Lookup dialog box appears.

Step 2 Select the **NS Lookup** tab.

Step 3 Enter an address or hostname in IPv4Address/Hostname.

Alternatively, click **Select** to choose a host network/host object that defines the IP address.

Step 4 Optionally, to employ a particular DNS server in the lookup, enter the server's name or address in DNS Server.

Step 5 Click **Lookup**.

The system displays the particular address/hostname pair, as well as the DNS server used in the lookup.

Using the Packet Capture Wizard

You can use the Packet Capture Wizard to configure, run, view, and save captures for troubleshooting errors. The captures can be run using preconfigured access lists or using match criteria of packet parameters such as source and destination addresses/ports on one or more interfaces. The wizard runs one capture on each of the ingress and egress interfaces. You can save the captures on the Cisco Security Manager client computer to examine them using a packet analyzer.

The Packet Capture Wizard also supports packet captures on ASA clusters. If you run the Packet Capture Wizard on the control unit of an ASA cluster, you are given the option of capturing data for just the selected device or all devices in the cluster. After running the capture for a cluster, you can view summary information for the cluster and also view or download capture buffers for specific devices in the cluster.

Note If the control unit has changed, it should be updated in Security Manager before running the Packet Capture Wizard. If not, capture for cluster members will contain errors. You can update the control unit for a cluster using the **Retrieve From Device** button on the **Device Properties > Cluster Information** page. For more information, see Cluster Information Page, page 3-50.

The captures can be run using pre-configured access-lists or using match criteria of packet parameters such as source, destination address/port on one or more interfaces.

Please note the following:

- You can only use the Packet Capture Wizard on firewall devices (PIX, ASA, or FWSM).
- Packet capture based on packet match criteria is only supported on devices running ASA version 7.2(3) or later. For other devices, packet capture can only be performed based on access-lists.

To use the Packet Capture Wizard:

Step 1 Launch the Packet Capture Wizard using one of the following methods:

- Select **Tools > Packet Capture Wizard**.
- (Device view) Right-click on an ASA, PIX, or FWSM device and select **Packet Capture** on the shortcut menu. Proceed to Step 3.
- (Event Viewer) Right-click on an event from an ASA, PIX, or FWSM device and select **Packet Capture** on the shortcut menu. Proceed to Step 3.

Step 2 If you launched the Packet Capture Wizard from the Tools menu, select the device on which you want to capture packets. The Security Devices list contains only devices on which packet capture can be run.

Step 3 If you selected a device that is the control unit of an ASA cluster, specify whether to run the capture for the selected device only or for the entire cluster, and then click **Next**.

Step 4 Select the ingress interface from the drop-down list.

Note You cannot select the same interface as ingress and egress in the same wizard.

Step 5 In the Packet Match Criteria area, do one of the following:

- To specify the access list to use for matching packets, select the **Access-List** radio button, and then choose the access list from the drop-down list.
- To specify packet parameters, select the **Packet Parameters** radio button and complete the following fields:
  - Specify the source and destination in the Source Host / Network and Destination Host / Network fields, respectively. You can use any of the following to specify the source or destination:
  - Source Host/ Network object. Enter the name of the object or click **Select** to select it from a list. You can also create new network/host objects from the selection list.

Note Starting from Cisco Security Manager 4.18, the packet parameters are supported with All-Address(any), All-IPv4-Address(any4), and All-IPv6-Address(any6).

  - Host IP address, for example, 10.10.10.100.
  - Network address, including subnet mask, in either the format 10.10.10.0/24 or 10.10.10.0/255.255.255.0.
  - Choose the protocol type to capture from the drop-down list. Available protocol types to capture are ah, eigrp, esp, gre, icmp, icmp6, igmp, igrp, ip, ipinip, nos, ospf, pcp, pim, snp, tcp, or udp.

    If the protocol is ICMP, select the ICMP type from the drop-down list. Available types include ALL, alternate-address, conversion-error, echo, echo-reply, information-reply, information-request, mask-reply, mask-request, mobile-redirect, parameter-problem, redirect, router-advertisement, router-solicitation, source-quench, time-exceeded, timestamp-reply, timestamp-request, traceroute, or unreachable.

    If the protocol is TCP or UDP, specify the source and destination port services. Available options include the following:
  - To include all services, choose All Services.

– To indicate specific services, choose an appropriate operator from the drop-down list (=, !=, >, <, or range) and then select one of the following: aol, bgp, chargen, cifs, citrix-ica, ctiqbe, daytime, discard, domain, echo, exec, finger, ftp, ftp-data, gopher, h323, hostname, http, https, ident, imap4, irc, kerberos, klogin, kshell, ldap, ldaps, login, lotusnotes, lpd, netbios-ssn, nfs, nntp, pcanywhere-data, pim-auto-rp, pop2, pop3, pptp, rsh, rtsp, sip, smtp, sqlnet, ssh, sunrpc, tacacs, talk, telnet, uucp, whois, or www. The >, <, and range operators function based on the port number assigned to the selected service.

  When using the range operator, a second drop-down list is enabled. Use the two drop-down lists to select the starting and ending services in the range you want to specify. The service with the lower corresponding port number should be selected in the first drop-down list, and the service with higher corresponding port number should be selected in the second drop-down list.

Step 6 Click **Next** to proceed to the Select egress interface step.

Step 7 Select the egress interface from the drop-down list.

Note You cannot select the same interface as ingress and egress in the same wizard.

Step 8 In the Packet Match Criteria area, do one of the following:

Note The Packet Match Criteria option (Access-list or Packet Parameters) you selected for the ingress interface will also be used for the egress interface. Also, if you used packet parameters for matching on the ingress interface, the protocol definition you used will also be used for the egress interface.

- If you are using an access list to match packets, choose the access list from the drop-down list.
- If you are using packet parameters to match packets, the parameters used for ingress are also used for egress.

Step 9 Click **Next** to proceed to the Set buffer parameters step.

Step 10 Specify the buffer parameters by configuring the following fields:

In the Buffer Parameters area, you specify the buffer size and packet size. The buffer size is the maximum amount of memory that the capture can use to store packets. The packet size is the longest packet that the capture can hold. We recommend that you use the longest packet size to capture as much information as possible.

- **Read capture buffer every 10 seconds**—Select this option to automatically retrieve captured data every 10 seconds. You must use the circular buffer when selecting this option.
- **Use a circular buffer**—Select this option to continue capturing packets after the buffer is full. When you choose this setting, if all the buffer storage is used, the capture starts overwriting the oldest packets.
- **Buffer Size**—Enter the number of bytes (between 1534 and 33554432) that the capture can use to store packets.
- **Maximum Packet Size**—Enter the number of bytes (between 14 and 1522) that the capture can use to store a single packet. Use the largest value, 1522,to capture as much information as possible.

Step 11 Click **Next** to proceed to the Summary step, which shows the traffic selectors and buffer parameters that you have entered.

Step 12 Click **Next** to proceed to the Run, View & Save step.

Step 13 From the Run, View & Save step, you can do the following:

- Click **Start Capture** to begin capturing packets.
- Click **Stop Capture** to stop capturing packets.
- To fetch the next set of captured packets, do one of the following:
  - For individual devices, click **Display Capture Packets** to fetch the next set of captured packets from the device and update the buffer status bar. This button is only enabled if the Read capture buffer every 10 seconds option was not selected during the Set buffer parameters step.
  - For clusters, click **Get Cluster Capture Summary** to fetch the next set of captured packets from the devices in the cluster and update the buffer status bar. This button is only enabled if the Read capture buffer every 10 seconds option was not selected during the Set buffer parameters step.
- When running a capture for an ASA cluster, the following option are available for working with the capture buffers of the devices in the cluster:
  - To view the captured packets from a device in the cluster, select the device in the Device Name list under Get Capture Buffer, and then click **Get Capture Buffer**.

    The capture information for the selected device is displayed. Refer to the other options in this list for the actions you can perform on this data.
  - To remove the capture content for a specific device or all devices in the cluster and allow room in the buffer to capture more packets, select the device or **--All--** in the Device Name field under Clear Capture Buffer, and then click **Clear Capture Buffer**.

    Note We recommend saving captures prior to clearing the device buffers. If you do not save captures prior to clearing the device buffers, captured data will be lost.
- Click the **Launch Network Sniffer** button above the Ingress Capture window or the Egress Capture window to view the corresponding ingress or egress capture using an external packet analyzer tool. You must have a packet analyzer installed and associated with *.pcap file extension.
- Click **View Data in Larger Window** to view the packet capture data side-by-side in a larger window.
- Click **Save captures** to display the Save Capture dialog box. Choose the format in which you want to include the captures: ASCII or PCAP. You have the option of saving either the ingress capture or the egress capture.
- Click **Clear Device Buffer** to remove the current content and allow room in the buffer to capture more packets.

  Note We recommend saving captures prior to clearing the device buffers. If you do not save captures prior to clearing the device buffers, captured data will be lost.
- Click **Refresh Capture Buffers** to fetch the next set of captured packets for a device in a cluster and update the buffer status bar.

Step 14 Click **Finish** to exit the wizard.

IP Intelligence

Network security devices managed by Cisco Security Manager generate large amounts of security logs and security events containing the IP address information of the attacker or victim machines or both.

Useful details about an IP address, collectively referred to as IP intelligence, can be discovered by using tools such as ping, trace route, and NS lookup. Often, however, it is desirable to augment these somewhat rudimentary tools with more advanced tools.

Beginning with Version 4.5, Security Manager provides advanced tools that furnish critical details about an IP address in real time or in generated reports. These critical details are provided by Security Manager in the following categories:

- Reverse DNS (FQDN) Lookup Service
- GeoIP Lookup Service
- Whois Lookup Service

Note IP intelligence for IPv6 addresses is not supported.

These IP intelligence categories are described in the following table:

Table 72-10 IP Intelligence Categories

IP Lookup Provider	Information Source	Real Time or Manual/Limitations
Reverse DNS (FQDN) Lookup Service	DNS servers	Real time **Note** External DNS configuration is an additional option that you can configure, but you will need to evaluate your individual situation.
GeoIP Lookup Service	External third-party commercial vendor	Real time

Table 72-10 IP Intelligence Categories (continued)

Whois Lookup Service	Provided by free whois server, a third-party web server.	Real time Limitations: • Whois is a query and response protocol that is widely used for querying databases that store the registered users or assignees of an Internet resource, such as a domain name, an IP address block, or an autonomous system. There are five regional internet registry (RIR) organizations that manage the allocation and registration of IP addresses. • ARIN (American Registry for Internet Numbers), RIPE (Réseaux IP Européens Network Coordination Centre), and APNIC (Asia-Pacific Network Information Centre) are the 3 RIR's that Security Manager will use to query directly; they also provide the referred URL. For RIPE and APNIC, if there is any parsing error, only the direct URL link will be shown. • Clicking on the provided URL will display the details for the given IP address in web browsers. If an IP address belongs to LACNIC (Latin America and Caribbean Network Information Centre) or AfriNIC (African Network Information Centre) , the web browser will display the homepage of the respective RIR. • In some cases (such as dns query getting blocked by windows firewall or invalid proxy has configured in the cco settings page) , Whois may not work even though it is enabled. In such scenario as a "Fail-safe" method it will provide only the referred url.

Before you begin looking up IP intelligence, you must enable the necessary services at **Configuration Manager > Tools > Security Manager Administration... > IP Intelligence Settings**. Refer to IP Intelligence Settings Page, page 11-41.

IP intelligence lookup can be done by using any of the following methods:

- Use the IP Intelligence dialog box: Navigate to **Configuration Manager > Tools > IP Intelligence...** and then enter a valid IPv4 address in the search field of the IP Intelligence dialog box that pops up. (You must press **Enter** after typing the IP address.)
- Use "Quick Launch" by mousing over (hovering over) a valid IPv4 address in the Security Manager interface. You can do this in Event Viewer, for example; in general, you can do this in all GUI tables in which an IP address is part of the data displayed. If more than one IP address is displayed in one cell of a GUI table, only the first IP address is displayed.

Note You may experience a second or two of latency with Quick Launch before you see the "IP Intelligence" option in the GUI.

Note You can enable or disable Quick Launch by selecting or clearing the "Enable Quick Launch" check box in the IP Intelligence dialog box at **Configuration Manager > Tools > IP Intelligence...**.

- Use the IP Intelligence widget in the Dashboard (**Launch > Dashboard...**). This method is equivalent to using the IP Intelligence dialog box described above.
- Use Report Manager (**Launch > Report Manager...**) to see IP intelligence in any of the following reports:
  - **FW/Summary Botnet**—Top Infected Hosts
  - **FW/Summary Botnet**—Top Malware Sites
  - **FW**—Top Destinations
  - **FW**—Top Sources
  - **IPS**—Target Analysis
  - **IPS**—Top Attackers
  - **IPS**—Top Victims

Note There are several points to note about these reports: 1) They will not contain Whois information. 2) If you have disabled all the providers at **Configuration Manager > Tools > Security Manager Administration... > IP Intelligence Settings**, then none of the IP intelligence-related columns will be displayed in the report. 3) If you enable *all* the services, only the reverse dns (FQDN) and GeoIP details will be displayed in the report. 4) If you enable only *one* service, then that service only will be displayed in the report.

Note The generated report in both PDF and CSV formats will contain the IP intelligence details.

Note All these reports require the necessary services to be enabled. Refer to IP Intelligence Settings Page, page 11-41.

Integrating CS-MARS and Security Manager

While Cisco Security Manager lets you centrally manage security policies and device settings in your network, the Cisco Security Monitoring, Analysis and Response System (CS-MARS) is a separate application that monitors devices and collects event information, including syslog messages and NetFlow traffic records, with much more extensive network monitoring capabilities than Security Manager. CS-MARS aggregates and presents massive amounts of network and security data in an easy-to-use format. Based on information derived from CS-MARS reports, you can edit device policies in Security Manager to counter security threats.

Specifically, if you use Security Manager to configure firewall access rules and IPS signatures, you can configure CS-MARS to collect information related to those policies and make it available to Security Manager users. By registering the CS-MARS servers with Security Manager, users can navigate directly from a specific access rule or IPS signature to a CS-MARS report window, pre-populated with query criteria for that rule or signature.

Similarly, CS-MARS users can view the Security Manager policies related to specific CS-MARS events. This bi-directional mapping of specific events to the policies that triggered them, combined with the ability to immediately modify the policies, can dramatically reduce the time spent configuring and troubleshooting large or complex networks.

To enable this cross-communication, you must register your CS-MARS servers with Security Manager, and register your Security Manager server with the CS-MARS servers. You must also register the specific devices with each application. Then, when working with firewall access rules or IPS signatures for a device, a Security Manager user can quickly view real-time and historical event information related to that rule or signature.

The following sections explain how to enable and use CS-MARS and Security Manager cross-communication:

- Checklist for Integrating CS-MARS with Security Manager, page 72-37
- Looking Up CS-MARS Events for a Security Manager Policy, page 72-42
- Looking Up a Security Manager Policy from a CS-MARS Event, page 72-46

Checklist for Integrating CS-MARS with Security Manager

To enable the cross-communication between CS-MARS and Security Manager (as described in Integrating CS-MARS and Security Manager, page 72-36), you must identify the applications to each other and ensure that devices managed by both applications are configured appropriately. The following table describes the integration steps.

If you have problems with cross-communications, see Troubleshooting Tips for CS-MARS Querying, page 72-41.

Table 72-11 Integrating CS-MARS and Security Manager

Task	Description
Add the devices to Security Manager and CS-MARS	See Adding Devices to the Device Inventory, page 3-7 for information about adding devices to Security Manager. See the *Device Configuration Guide for Cisco Security MARS* for information about adding devices to the CS-MARS inventory. A device must be supported by both applications to provide cross-communication for the device. Supported device types generally are those providing Firewall > Access Rules, or IPS > Signatures policies. (These include: PIX, ASA and FWSM appliances, Cisco IOS routers, Cisco IPS sensors and modules, and Cisco Catalyst switches.)
Configure the devices as required by each application	See Chapter 2, "Preparing Devices for Management" for information about basic configuration requirements for Security Manager. See *Device Configuration Guide for Cisco Security MARS* for the more extensive requirements for CS-MARS.
Register Security Manager with CS-MARS	For information on configuring CS-MARS to communicate with Security Manager, see *User Guide for Cisco Security MARS Local and Global Controllers*. You might want to create a CS-MARS user account specifically for linking with Security Manager. See Configuring the Security Manager Server to Respond to CS-MARS Policy Queries, page 72-38.
Register CS-MARS controllers with Security Manager	For information on registering CS-MARS controllers with Security Manager, see Registering CS-MARS Servers in Security Manager, page 72-39.
Link CS-MARS controllers to the devices in Security Manager	In Security Manager, you can proactively discover the CS-MARS controllers that monitor a particular device by clicking **Discover CS-MARS** on the device's Device Properties page, as described in Discovering or Changing the CS-MARS Controllers for a Device, page 72-40. Otherwise, the appropriate controller is discovered automatically when a user attempts to look up events for the device (the user is prompted to select a controller if more than one monitors the device).

Related Topics

- Viewing CS-MARS Events for an Access Rule, page 72-43
- Viewing CS-MARS Events for an IPS Signature, page 72-45
- Looking Up a Security Manager Policy from a CS-MARS Event, page 72-46

Configuring the Security Manager Server to Respond to CS-MARS Policy Queries

CS-MARS must be allowed access to the Security Manager server so that it can perform policy lookup queries and obtain policy information.

- If you are using Common Services AAA authentication on the server (for example, Cisco Secure ACS), you must update the administrative access settings to ensure that CS-MARS has the necessary client access to the Security Manager server.

- Define a user account in Security Manager that CS-MARS can use to perform queries. A separate account is recommended to provide a specific audit trail on the Security Manager server. This account must be assigned one of the following Common Services roles:
  - Approver
  - Network Operator
  - Network Administrator
  - System Administrator

Users with the Help Desk security level can only view the policy look-up table in CS-MARS; that is, they cannot cross-launch Security Manager to modify policies.

Note When you register a Security Manager server with CS-MARS, if you choose to prompt for Security Manager credentials for policy table look-up, a separate CS-MARS account in Common Services for authentication purposes might not be necessary.

For more information on adding users and associating roles with them in Common Services, see the *User Guide for CiscoWorks Common Services*.

Related Topics

- Registering CS-MARS Servers in Security Manager, page 72-39
- Discovering or Changing the CS-MARS Controllers for a Device, page 72-40

Registering CS-MARS Servers in Security Manager

As described in Checklist for Integrating CS-MARS with Security Manager, page 72-37, you must register your CS-MARS controllers with Security Manager to enable cross-communication between the applications if you intend to use the applications together.

Then, when a user looks up events for a device, Security Manager identifies the CS-MARS controller that is collecting events for that device. If more than one CS-MARS controller is collecting events for a device, the user can select which to use. You can also specify the correct CS-MARS controller to use in the Device Properties window for each device. (See Discovering or Changing the CS-MARS Controllers for a Device, page 72-40 for more information.)

Note For information about the CS-MARS versions explicitly supported by Security Manager, see the *Release Notes for Cisco Security Manager* for this version of the product. If you do try to use a version that is not explicitly supported, you cannot use CS-MARS versions earlier than 4.3.4 or 5.3.4.

Step 1 Choose **Tools > Security Manager Administration** and select **CS-MARS** in the table of contents to display the CS-MARS Page, page 11-7.

Step 2 Click the **Add** button to add a CS-MARS server. The New CS-MARS Device dialog box opens (see New or Edit CS-MARS Device Dialog Box, page 11-8 for detailed information).

Step 3 In the New CS-MARS Device dialog box, enter the IP address or fully qualified DNS host name of the server, and a user name and password for logging into the server. If you add a local controller, the user name you enter can be either a local account or a global account. Choose the type of account from the User Type list.

Tip If you are using CS-MARS Global Controllers, add them instead of individual Local Controllers. By adding Global Controllers, Security Manager can identify the correct Local Controller for a device, without you having to add each Local Controller. When you add a Global Controller, do not add the individual Local Controllers monitored by the Global Controller.

Click **Retrieve From Device** to get the server's authentication certificate. Click **Accept** when the certificate is presented to you.

Click **OK** when finished. The New CS-MARS Device dialog box closes and the server is added to the CS-MARS device list.

Step 4 From the **When Launching CS-MARS** list, choose whether you want users to be prompted to log in to the CS-MARS server when they request event status, or whether Security Manager should automatically log in to CS-MARS using the credentials provided when the user logged in to Security Manager.

If you elect to use Security Manager credentials, the necessary user accounts must be configured in CS-MARS. Refer to the CS-MARS documentation for more information.

Step 5 Click **Save** on the CS-MARS page to save your changes.

Discovering or Changing the CS-MARS Controllers for a Device

If you use the Cisco Security Monitoring, Analysis and Response System (CS-MARS) controllers to monitor devices, you can register them in Security Manager and then view syslogs and events that are related to firewall access or IPS signature rules for individual devices.

Security Manager can automatically discover the CS-MARS controllers that monitor a device when you try to view events related to a rule. If more than one controller monitors a device, you are prompted to select which controller to use.

You can also proactively select the CS-MARS controller for a device in its Device Properties window. Similarly, if you ever need to change the CS-MARS controller assigned to a device, you can change the selection in its Device Properties window. This procedure explains how to discover or change the CS-MARS controller for a device from its Device Properties window.

Before You Begin

The CS-MARS controller that monitors the device must already be registered with Security Manager on the CS-MARS administration page (**Tools > Security Manager Administration > CS-MARS**). For more information, see Registering CS-MARS Servers in Security Manager, page 72-39.

Step 1 In Device view, do one of the following in the Device selector to open the Device Properties dialog box:

- Double-click a device.
- Right-click a device and choose **Device Properties**.
- Select a device and choose **Tools > Device Properties**.

Step 2 Click **General** in the table of contents to open the General properties page (see Device Properties: General Page, page 3-41).

Step 3 In the CS-MARS Monitoring group, click **Discover CS-MARS**. Security Manager determines which registered controller is monitoring the device, if any. If there are more than one, you are prompted to select which CS-MARS controller to use.

Troubleshooting Tips for CS-MARS Querying

Use the following troubleshooting tips to help you identify and resolve problems you might encounter when using CS-MARS and Security Manager together:

- HTTPS is required for communication between the Security Manager server and CS-MARS.
- Interface names are not case-sensitive in Security Manager, but they are in CS-MARS. For example, "outside" and "Outside" are considered exclusive by a CS-MARS appliance, while they are equivalent in Security Manager. Further, syslog messages use lower case for all interface names. As a result, when you perform a query for a Security Manager policy from an event generated in CS-MARS, the interface name logged in the syslog event might not match the interface name in that policy in Security Manager. To avoid this problem, use lower case for all interface names, and in the definition of interface roles, in CS-MARS.
- To query for CS-MARS events from Security Manager policies, the Security Manager client must be on the same side of a network address translation (NAT) boundary as the CS-MARS appliance and the Security Manager server.

  Similarly, when the CS-MARS client is not on the same side of a NAT boundary as the CS-MARS appliance and the Security Manager server, you can look up Security Manager policies, but in read-only mode. However, you cannot start the Security Manager client from the read-only policy look-up table. The Security Manager client must be on the same side of the NAT boundary as the CS-MARS appliance and the Security Manager server if you want to start the client from CS-MARS to modify a matching policy.
- For FWSM, PIX and ASA devices on which multiple independent security contexts exist, to query for CS-MARS events, you must define a unique management IP address in Security Manager for each security context. Also, the host name and reporting IP address for each virtual context must be configured before adding it to CS-MARS. Otherwise, event look-up from policies on these contexts fails.
- For all IPS device and service policies, a default signature policy is assigned to the device when you do not discover IPS policies, or when you remove the configured policies from the device. If you try to perform event look-up from the default signature, a "Policy not found" error message is displayed. However, if you edit the default signature and save it, you can then navigate to events in CS-MARS.
- If object grouping or rule optimization is enabled for an access rule defined in Security Manager and the associated access-list commands on the device do not match the optimized rules, no events are displayed in CS-MARS.
- If logging is not enabled for an access rule, a warning message is displayed, and you can only look up traffic-flow events for those rules.
- When supported by the device, Security Manager uses access-control entry (ACE) hashcodes as additional keywords when querying CS-MARS for syslog messages generated by an ACE, and large access-control lists (ACLs) might contain thousands of such hashcodes. If the number of keywords, or the sum of the number of sources, destinations, and protocols for an ACE or a signature exceeds the query limit of 150, an error message is displayed. The error message indicates the probable cause and recommended action.
- Problems with the synchronization between rules and reported events can occur in the following situations:

  - The device has been added to Security Manager, but the configuration or changes to it have not been saved to the database. This is especially true for access rules that have been changed but not deployed since the device was added to CS-MARS.
  - Access rules exist on the device for which there are no corresponding rules in Security Manager, and vice versa. Be sure all devices are added to Security Manager, and that access rules are configured on them using Security Manager.
  - Traffic in the "wrong" direction triggering events for which there is no defined rule. For example, outbound traffic on a higher-security-level interface on which only inbound-traffic rules have been defined.
- If you perform a policy lookup from CS-MARS and the Security Manager client is active, the query is performed on all policies within the open activity or configuration session plus what is saved in the database (the committed configurations). If the Security Manager client is not active, only committed policies are considered.

Related Topics

- Checklist for Integrating CS-MARS with Security Manager, page 72-37
- Looking Up CS-MARS Events for a Security Manager Policy, page 72-42
- Registering CS-MARS Servers in Security Manager, page 72-39

Looking Up CS-MARS Events for a Security Manager Policy

After you integrate CS-MARS and Security Manager, you can look up events in CS-MARS that relate to specific firewall access rules or IPS signatures.

When CS-MARS receives events, they are parsed, "sessionized," written to an event buffer, and then written to the database. Sessionizing takes two forms: with a session-oriented protocol, such as TCP, the session encompasses the initial handshake to the connection tear-down; with a sessionless protocol, such as UDP, the session start and end times are based more on first and last packets tracked within a restricted time period—packets that fall outside of the time period are considered parts of other sessions.

Because there is a difference between newly-received and fully processed data, you can look up either real-time or historical events:

- **Real-time**—Because sessionization takes time, keeping an event in cache for up to two minutes, you can use the real-time event query to view events right after parsing, providing access to the most current data received.

  When you query for real-time events, the query is run automatically, based on the policy values obtained from Security Manager, and the results are displayed in the CS-MARS Query Results window. This real-time event viewer lets you monitor CS-MARS traffic in near real-time, as raw events streaming to CS-MARS, before they are sessionized, with a maximum delay of five seconds. You also can elect to view the sessionized event stream by clicking Edit in the Query Results window and then choosing "Sessionized events" from the Realtime drop-down menu. Note that more delay is possible when there are many events in a session.
- **Historical**—Historical event reports help you identify trends over longer periods of time than is possible with real-time monitoring. When you query for historical events, the CS-MARS Query Criteria: Result window opens. You can either run the query immediately, or save the criteria as a "report" to run at a later time. For historical events, the Result Format is the All Matching Events option, and the Filter By Time value is set to the previous 10 minutes.

The following topics explain event lookup in more detail:

- Viewing CS-MARS Events for an Access Rule, page 72-43
- Viewing CS-MARS Events for an IPS Signature, page 72-45

Viewing CS-MARS Events for an Access Rule

From the **Firewall > Access Rules** policy in Security Manager, you can select an access rule and view related event information in CS-MARS. You can view real-time or historical events matching the rule, the traffic flow, the source address, or the destination address. You can view events for any device that supports access rules, including ASA, PIX, FWSM, routers, and switches.

Firewall access rules are presented in the form of an ordered list or table. When deployed, this policy becomes an access-control list (ACL), with each entry in the list known as an access-control entry (ACE). (For more detailed information, see Understanding Access Rules, page 16-1.)

When deciding whether to forward or drop a packet, a device tests the packet against each access rule in the ordered listed. If you enable logging for an access rule, the results of the test are recording according to your per-rule log settings. Some devices, such as ASA, generate log entries for denied access even if you do not configure logging explicitly. For information on creating access rules, including logging options, see Configuring Access Rules, page 16-7.

You can query CS-MARS for real-time or historical events related to an access rule for the following types of traffic. To use the commands, right-click the rule and select them from the context menu.

- **Flow**—A traffic flow is defined by the rule's source and destination IP addresses, protocol, and ports. The reported flow events include connection set-up and tear-down. Logging need not be enabled for the access rule to record this information.

  To view flow-related events, use the following right-click commands:

  - **Show MARS Events > Realtime > Matching this Flow**—To view real-time query results in CS-MARS for events matching this traffic flow. You can change the query criteria in the CS-MARS window at any time, applying new parameters to alter the real-time results.
  - **Show MARS Events > Historical > Matching this Flow**—Opens the historical query criteria page in CS-MARS with fields populated based on the selected rule's traffic flow. Edit the rule parameters and query criteria as desired, and click **Apply** to continue. Next, in the Query window, you can submit the query or save it for later submission and re-use.

- **Rule**—If logging is enabled for the rule (in the Advanced and Edit Options Dialog Boxes, page 16-17), the device sends syslog messages to CS-MARS to record the logged events (assuming CS-MARS monitors the device). This query includes the access-rule parameters, including available keyword information. Reported events do not include connection set-up and tear-down.

  To view rule-related events, use the following right-click commands:

  - **Show MARS Events > Realtime > Matching this Rule**—To view real-time query results in CS-MARS for events matching this rule (flow parameters plus keywords); results begin scrolling within five seconds. You can change the query criteria in the CS-MARS window at any time, applying new parameters to alter the real-time results.
  - **Show MARS Events > Historical > Matching this Rule**—Opens the historical query criteria page in CS-MARS with fields populated based on the access rule (flow parameters plus keywords). Edit the rule parameters and query criteria as desired, and click **Apply** to continue. Next, in the Query window, you can submit the query, or save it for later submission and re-use.

- **Source or Destination**—If you right-click the Source or Destination cell in an access rule entry, you also can choose to view real-time or historical events matching the rule's source or destination IP address.

To view events for a source or destination address, right-click the address in the Source or Destination cell and choose one of the following commands (the specific command differs depending on the cell you select):

- **Show MARS Events > Realtime > Matching this Source/Destination**—To view real-time query results in CS-MARS for events with a matching source or destination address. You can change the query criteria in the CS-MARS window at any time, applying new parameters to alter the real-time results.
- **Show MARS Events > Historical > Matching this Source/Destination**—Opens the historical query criteria page in CS-MARS with fields populated based on the access rule's source or destination address. Edit the rule parameters and query criteria as desired, and click **Apply** to continue. Next, in the Query window, you can submit the query, or save it for later submission and re-use.

Security Manager provides the following information to CS-MARS as criteria for a traffic-flow or access-rule event queries:

- Device details—General information about the device, such as host name, domain name, management IP address, and display name.
- Source addresses—Source addresses of hosts and the network/host objects expanded to display the networks or collections of IP addresses.
- Destination addresses—Destination addresses of hosts and the network/host objects expanded to display the networks or collections of IP addresses.
- Service—Protocol and port information.
- Event Type—"Built/teardown/permitted IP connection" for permit rules and "Deny packet due to security policy" for deny rules.
- Keyword (rule events only, not provided for traffic-flow queries)—ACL name and ACE hashcode, if available, connected by the logical operator OR.

  On Version 7.0 or later PIX and ASA devices, each access rule is assigned an MD5 hashcode, which is included in the syslogs generated by that rule. Large ACLs can include thousands of access rules. Used as query keywords, these hashcodes can help produce more-accurate event matches. If a device does not support hashcodes, a warning is displayed that query results might be inaccurate because of keyword ambiguity; you can proceed with the query, and then edit the query keyword list and resubmit.

Tips:

- You can query on only one access rule at a time.
- When NAT or PAT is configured on a security device, the source and destination addresses are mapped to pre-translation and post-translation addresses, respectively, and the translated addresses are used when Security Manager sends a query to CS-MARS. For inbound access rules, the destination address is considered the pre-translation address, and for outbound access rules, the source address is considered the post-translation address.
- If the device is monitored by multiple CS-MARS controllers, you are prompted to select the CS-MARS instance to be used.
- Depending on how credentials verification is set up on your system, you might be prompted to log into CS-MARS. For more information, see Registering CS-MARS Servers in Security Manager, page 72-39.

Related Topics

- Access Rules Page, page 16-10

- Looking Up CS-MARS Events for a Security Manager Policy, page 72-42
- Viewing CS-MARS Events for an IPS Signature, page 72-45

Viewing CS-MARS Events for an IPS Signature

Note From version 4.17, though Cisco Security Manager continues to support IPS features/functionality, it does not support any bug fixes or enhancements.

When an IPS or IOS IPS device detects and reports a network intrusion by comparing incoming traffic to a configured signature, a syslog message is generated on the device. If the device is monitored by CS-MARS, an incident is generated in CS-MARS after the log associated with the signature is obtained from the device. Looking up the events associated with a specific signature lets you quickly identify attacks and tune your device configuration to minimize or prevent intrusions.

To view reported network intrusion events in CS-MARS, you can select one or more entries in the Signatures policy for a device in Security Manager and navigate to the CS-MARS Query page to view real-time and historical events.

When you look up real-time events for a signature, the query is run automatically and the results displayed in CS-MARS. However, when you look up historical events for a signature, the values sent by Security Manager to CS-MARS are used to populate the query fields. You can modify the query fields as desired, and then run the query, or save it for later use.

Security Manager provides the following signature information to CS-MARS as query criteria:

- Device details—General information about the device, such as host name, domain name, management IP address, and display name.
- Keyword—Signature ID, subsignature ID, and virtual sensor name, if applicable.

  For virtual sensors, the name of the sensor is included as a keyword criterion along with other device information and signature parameters.

Related Topics

- Looking Up CS-MARS Events for a Security Manager Policy, page 72-42
- Viewing CS-MARS Events for an Access Rule, page 72-43

Step 1 (Device view) With an IPS or IOS IPS device selected, select **IPS > Signatures > Signatures** to display the Signatures Page, page 39-4.

Step 2 Right-click the desired entry in the signatures table, or select multiple entries before right-clicking one of them, and choose one of the following commands from the **Show MARS Events** menu:

- **Realtime**—To view real-time query results in CS-MARS for events matching this signature; results begin scrolling within five seconds. Use this option to view raw events as they stream to CS-MARS.

  You can change the query criteria in the CS-MARS Query Results window at any time, applying new parameters to alter the real-time results.
- **Historical**—Opens the historical query criteria page in CS-MARS with fields populated based on the signature parameters. Edit the parameters and query criteria as desired, and click Apply to continue. Next, in the Query window, you can submit the query or save it for later submission and re-use. You can edit the query and save it as a report if you want to run it again later.

Tips:

- If a signature is disabled, you are warned and asked if you want to proceed to event lookup.
- If the device is monitored by multiple CS-MARS controllers, you are prompted to select the CS-MARS instance to be used.
- Depending on how credentials verification is set up on your system, you might be prompted to log into CS-MARS. For more information, see Registering CS-MARS Servers in Security Manager, page 72-39.
- All custom signatures are categorized as "Unknown Device Event Type" events in CS-MARS.
- A default signature is assigned to an IPS device if you elect not to discover IPS policies when adding the device to the Security Manager inventory, or when you remove configured IPS policies from the device. If you try to look up events from the default signature, a "Policy not found" error message is displayed. However, if you edit the default signature and save it, you can then query for related events in CS-MARS.
- Events of type Packet Data and Context Data are not displayed in the query results because these events are not triggered by signature rules.

Looking Up a Security Manager Policy from a CS-MARS Event

The *User Guide for Cisco Security MARS Local and Global Controllers* contains detailed information about how to look up policies based on events shown in CS-MARS. The information includes extensive troubleshooting information to help resolve any problems you might have, plus a checklist of what you must configure in CS-MARS to enable the interaction.

The main reason you would want to perform policy lookup is to adjust a policy based on the events that it is generating. For example, an access rule might be dropping traffic that you actually want to allow. Because you are looking at the event, you know there is a policy that is causing the event, so with a few clicks, you can get from that event to the policy you need to reconfigure.

The general process for looking up a policy based on a device-generated event is as follows. Note that the Security Manager client must be installed on your system to perform policy lookup.

Related Topics

- Viewing CS-MARS Events for an Access Rule, page 72-43
- Viewing CS-MARS Events for an IPS Signature, page 72-45

Step 1 Find the event in CS-MARS in the Query Results or Incident Details pages.

For more information on the syslog and NetFlow events you can use for querying access rules, see the following topics:

- System Log Messages Supported for Policy Look-up, page 72-47
- NetFlow Event Reporting in CS-MARS, page 72-48

Step 2 Click the Security Manager icon in the Reporting Device cell for the event. You might be prompted to log into Security Manager, based on how you configured CS-MARS.

If more than one device in Security Manager matches the event characteristics, you are prompted to select a device.

Step 3 Detailed information is obtained from Security Manager and presented based on whether the event is for an access rule or IPS signature:

- **Access rule**—The access rules are displayed in CS-MARS in a read-only window with the rule that matches the event highlighted.

  If you decide to edit a rule, click the rule number, and you are taken to the rule in the Access Rule policy in the Security Manager client. You can then make your edits, save them, and then deploy configurations. Remember that your changes are not made to the device until you deploy them.

  For more information on configuring access rules, see Configuring Access Rules, page 16-7.

- **IPS Signature**—Signature details are displayed in CS-MARS in a read-only window.

  To edit the signature, click **Edit Signature**, and you are taken to the signature in the Signatures policy, where you can make your changes. For more information, see Editing Signature Parameters (Tuning Signatures), page 39-23.

  If you decide you want to instead remove specific actions from an event, or remove the event entirely, and prevent further processing by the sensor, click **Add Filter**. This opens the Add Event Filter dialog box in Security Manager, where you can configure an event filter. For more information, see Filter Item Dialog Box, page 40-9.

  As with access rules, your changes do not take effect until you deploy the new configuration.

System Log Messages Supported for Policy Look-up

When you configure access rules on security appliances and IOS devices, you can configure logging options in the Advanced and Edit Options Dialog Boxes, page 16-17 that generate system log (syslog) messages. On devices with multiple contexts, each security context includes its own logging configuration and generates its own messages. If Security Manager is configured to interoperate with CS-MARS, these messages are reported to CS-MARS and you can query for the reported information on a per-rule basis.

For additional information about each of these message IDs, see the System Message Guide of the relevant product documentation.

Security-appliance messages

Security-appliance syslog messages begin with a percent sign (%) and are structured as follows:

```
%{ASA | PIX | FWSM}-Level-Message_number: Message_text
```

For example:

```
%ASA-2-302013: Built outbound TCP connection 42210
for outside:9.1.154.12/23 (9.1.154.12/23) to inside:2.168.154.12/4402
(192.168.154.12/4402)
```

Note that additional information, such as date and timestamp, precedes these messages. The specific additional information depends on the type of device.

A unique six-digit number identifies each message (*302013* in the preceding example). The following security-appliance syslog message IDs are supported for Security Manager-to-CS-MARS queries. If you change the logging level of a security appliance, be sure these messages are generated at the new level.

Message ID	Message
106023	An IP packet was denied by the access rule. This message is recorded even if logging is not enabled for the rule; this is the Default Logging option.

106100	An IP packet was permitted or denied by the access rule. Additional information is provided, based on the logging level defined for the rule in the Advanced and Edit Options Dialog Boxes, page 16-17.
302013	A TCP connection between two hosts was created.
302014	A TCP connection between two hosts was torn down.
302015	A UDP connection between two hosts was created.
302016	A UDP connection between two hosts was torn down.
302020	A ICMP connection between two hosts was created.
302021	A ICMP connection between two hosts was torn down.

Router messages

On Cisco IOS routers, syslog messages are also generated for access rules. The first packet that triggers the access list causes an immediate logging message, and subsequent packets are collected over five-minute intervals before they are displayed or logged. Each logging message includes the access list number, whether the packet was permitted or denied, the source IP address of the packet, and the number of packets from that source permitted or denied in the prior five-minute interval.

The following IOS syslog message IDs are supported for Security Manager-to-CS-MARS queries.

%SEC-6-IPACCESSLOGP	A packet matching the log criteria for the given access list was detected: TCP and UDP.
%SEC-6-IPACCESSLOGS	A packet matching the log criteria for the given access list was detected: IP address.
%SEC-6-IPACCESSLOGDP	A packet matching the log criteria for the given access list was detected: ICMP.
%SEC-6-IPACCESSLOGNP	A packet matching the log criteria for the given access list was detected: all other IPv4 protocols.

Note If an excessive number of syslog messages are being generated and reported to CS-MARS, use the Advanced and Edit Options Dialog Boxes, page 16-17 to change the logging level for those access rules that are producing the largest number of messages. You can also look at changing the logging policies on the device to limit the types of messages generated.

NetFlow Event Reporting in CS-MARS

Event reporting in CS-MARS can include NetFlow events from an ASA 8.1+ device.

NetFlow Security Event Logging uses NetFlow version 9 fields and templates to efficiently deliver security telemetry in high-performance environments. NetFlow Security Event Logging scales better than syslog messaging, while offering the same level of detail and granularity in logged events. The ASA NetFlow implementation exports only significant events in the life of a flow, rather than exporting data about flows at regular intervals. The following flow events are exported:

- Flow creation
- Flow tear-down
- Flows denied by an access rule

The ASA also exports syslog messages that contain the same information. If you enable NetFlow on a device, you can consider disabling the equivalent syslog messages. Disabling equivalent syslog messages can help avoid the potential performance degradation caused by generating and processing both NetFlow records and syslog messages representing the same event. The following table lists syslog messages with an equivalent NetFlow event; the NetFlow Event IDs and Extended Event IDs are included. For information on how to disable NetFlow equivalent syslog messages, see Server Setup Page, page 54-21.

Syslog ID	Syslog Description	NetFlow Event ID	Extended Event ID
302013 302015 302017 302020	TCP, UDP, GRE, and ICMP connection creation.	1 = Flow Created.	0 = Ignore.
302014 302016 302018 302021	TCP, UDP, GRE, and ICMP connection tear-down.	2 = Flow Deleted.	0 = Ignore, or > 2000 = ASP drop reasons.
710003	An attempt to connect to the device's interface was denied.	3 = Flow Denied.	1003 = To-the-box flow denied due to configuration.
106015	A TCP flow was denied because the first packet was not a SYN packet.	3 = Flow Denied.	1004 = Flow denied because first packet was not a TCP SYN packet.
313001	An ICMP packet to the device was denied.	3 = Flow Denied.	1003 = To-the-box flow denied due to configuration.
313008	An ICMP v6 packet to the device was denied.	3 = Flow Denied.	1003 = To-the-box flow denied due to configuration.
106023	A flow was denied by an access list attached to an interface with the access group command.	3 = Flow Denied.	1001 – Flow denied by Ingress ACL. 1002 – Flow denied by Egress ACL.
106100	An access rule was hit.	1 = Flow Created (if ACL permitted the flow). 3 = Flow Denied (if ACL denied the flow).	0 – If Flow permitted by ACL. 1001 – Flow denied by Ingress ACL. 1002 – Flow denied by Egress ACL.

For the Flow Denied NetFlow event, an Extended Event ID indicates the reason for denial, as shown in the following table.

Extended Event ID	Event	Description
1001	FLOW DENIED	The flow was denied by an Ingress ACL.
1002	FLOW DENIED	The flow was denied by an Egress ACL.

Extended Event ID	Event	Description
1003	FLOW DENIED	The security appliance denied an attempt to connect to the interface service. For example, this message appears (with the service SNMP) when the security appliance receives an SNMP request from an unauthorized SNMP management station.
1004	FLOW DENIED	The flow was denied because the first packet was not a TCP SYN packet.
> 2000	FLOW DELETED	Values above 2000 represent various reasons for a flow being terminated.

PART 8

Image Management

Using Image Manager

Image Manager is a tool to simplify the distribution and management of images on internal and edge firewall devices in your network. It enables you to:

- Download and maintain a repository of different types and versions of images
- Evaluate images
- Analyze impact of upgrading images to the devices in the network
- Prepare for and plan an upgrade
- Reliably upgrade devices, with sufficient fallback and recovery mechanisms built in to ensure minimal network downtime

This chapter contains the following topics:

- Getting Started with Image Manager, page 73-1
- Working with Images, page 73-9
- Working with Bundles, page 73-13
- Working with Devices, page 73-16
- About Image Updates on Devices Using Image Manager, page 73-20
- Working with Jobs, page 73-32
- Troubleshooting Image Management, page 73-37

Getting Started with Image Manager

Image Manager contains sections that are used for managing your images, working with the devices that you will need to update, and performing image installations on those devices.

For more information on these areas of Image Manager, see the following topics:

- Working with Images, page 73-9
- Working with Bundles, page 73-13
- Working with Devices, page 73-16
- Working with Jobs, page 73-32

Before working with Image Manager, you should review the sections that follow:

- the platforms that are supported by this feature,
- the configuration settings you can change to control how the feature works, and

- the steps that are necessary to ensure your devices are configured to work with Image Manager.

This section contains the following topics:

- Image Manager Supported Platforms and Versions, page 73-2
- Device Configurations supported by Image Manager, page 73-4
- Image Manager Supported Image Types, page 73-5
- Administrative Settings for Image Manager, page 73-6
- Bootstrapping Devices for Image Manager, page 73-8

Image Manager Supported Platforms and Versions

Caution From version 4.18, Cisco Security Manager does not support SFR from ASA 9.10(1) onwards for ASA 5512, ASA 5506, ASA 5506H and ASA 5506W models. Therefore, if you upgrade to 9.10(1) through Image Manager, the exiting SFR configuration will be lost.

Image Manager is available only for ASA devices. The following devices support Image Manager:

- All legacy ASA models—ASA 5505/10/20/40/50/80
- ASA 5585
- ASA 5515/25/35/45/55
- ASA-SM module for Catalyst 6K
- 5516-X
- Adaptive Security Virtual Appliance (ASAv)

The following devices are not supported and are filtered out in the devices tab of the Image Manager unified view:

- PIX firewall
- FWSM blade
- ASA device managed by AUS
- Devices unmanaged in Security Manager
- ASA devices running on Firepower platform
- Other device types—IPS and Routers

Image Manager supports image upgrade for ASA device version from 7.x onwards. The target image version that can be used to upgrade is not restricted. Image upgrade to the highest ASA version supported in Security Manager 4.4, that is ASA version 9.0(1) and 9.1(1), has been tested.

Prior to version 4.9, the Image Manager application listed all the images of the supported device type. You could select and download any image that you required. Beginning with version 4.9, the Image Manager application lists only the specific versions of images.

The latest images of ASDM, †Remote Access Plugin and Host scan are listed in Image Manager. For AnyConnect version 3.x and 4.x, the latest images are listed.

For ASA devices, the following images are listed:

ASA Device Model	ASA Images listed in Image Manager
5512-x,5515-x,5525-x,5545-x,5585x	9.4.1
	9.3.3
	9.3.2
	9.3.1.SMP
	9.2.3.SMP
	9.2.2.4.SMP
	9.2.1.SMP.ED
	9.1.4.SMP.ED
	9.1.5.SMP.ED
	9.1.6.SMP.ED
	9.1.2.SMP.ED
	9.0.4.SMP.ED
	8.4.6.SMP.ED
5580-x	9.1.6.SMP
	9.1.5.SMP.ED
	9.1.4.SMP.ED
	9.1.2.SMP.ED
	9.0.4.SMP.ED
	8.4.6.SMP.ED
5555-x	9.4.1
	9.3.3
	9.3.2
	9.3.1.SMP
	9.2.3.SMP
	9.2.2.4.SMP
	9.2.1.SMP.ED
	9.1.2.SMP.ED
	9.0.4.SMP.ED
5505,5510,5520,5540,5550	9.1.6
	9.1.5.ED
	9.1.4.ED
	9.1.2.ED
	9.0.4.ED
	8.4.6.ED
5506-X	9.4.1, 9.3.3, 9.3.2

5506H-X	9.4.1
5506W-X	9.4.1
5516-X	9.4.1
Adaptive Security Virtual Appliance (ASAv)	9.3.1, 9.3.2, 9.4.1

Warning **Image downgrade is not restricted, but is done at your own risk. Image Manager does not validate downgrade cases.**

Device Configurations supported by Image Manager

In addition to supporting image updates on standalone ASA devices, Image Manager manages the filesystem and supports seamless image update for ASA devices specially configured for high availability and scalability. Following configurations are supported:

- **Multiple context mode**—ASA in multiple context mode where a single ASA can be partitioned into multiple virtual devices/firewalls. Refer to *http://www.cisco.com/en/US/docs/security/asa/asa91/configuration/general/ha_contexts.html*. Each of these virtual firewalls are represented in Security Manager as independent devices. When Image Manager updates the image on the physical unit hosting these virtual devices, it updates device properties of all virtual devices with the new image information.
- **Failover configuration**—Two identical ASA devices configured to failover for high availability. They can be configured to be in Active/Active or Active/Standby failover. Refer to *http://www.cisco.com/en/US/docs/security/asa/asa91/configuration/general/ha_overview.html*. Image update on an Active/Active failover pair is not supported in Image Manager. In order to use Image Manager to update the images on an Active/Active failover pair, the Active/Active failover pair has to temporarily be converted to Active/Standby by making all the failover groups active on one unit, and the corresponding failover groups standby on the other unit. After upgrade, you can convert the failover pair back to Active/Active.
- **Cluster configuration**—Multiple ASAs (up to 8 ASAs) can be grouped together as a single logical unit called a **cluster** for achieving increased throughput and redundancy. The purpose of clustering devices is to simplify manageability and to increase processing speed. By using clusters you are able to scale to a multitude of simultaneous connections that work together to load balance the connections. Clustering feature has been introduced starting from ASA version 9.0(1). For more information see *http://www.cisco.com/en/US/docs/security/asa/asa91/configuration/general/ha_cluster.html*.

Note Clustering is only supported on ASA 5580 and ASA 5585.

Starting from release 4.4, Security Manager supports Clustering. In Configuration manager and Image Manager, all the devices/members in a cluster or a failover pair are managed as a single device. That is, when you change the configuration on a cluster's control unit, the change is automatically made to all the devices in the cluster. Similarly, Image Manager updates image on each of the physical unit that is part of failover or cluster in a single operation.

Image Management for Multi-Context ASA

Beginning with version 4.12, the Image Manager Device Tree view displays all the user contexts (Admin and User contexts) of the multi-context firewall devices running the ASA software version 9.6(2) or later.

You can select a User context and view the storage-url information of the selected context on the Storage tab.

On the Compatible Images tab, you can view only the AnyConnect images for the selected User Context. However, all image types are displayed for the System Contexts.

Image Manager Supported Image Types

Image Manager supports the following types of images:

- ASA System software
- ASDM image
- VPN images [includes Cisco Secure Desktop (CSD), AnyConnect, and Hostscan]
- SSLVPN Plug-in images (For example: RDP, SSH, ICA, and others)

Image Manager completely manages the ASA system software and the ASDM images on the ASA devices, i.e., it performs loading of the image, activating the image by modifying configuration, and even reloading the device if required to complete the image upgrade process.

For the User Context devices Security Manager supports only AnyConnect images for copying and installation.

Image Manager does not support the ASA–CX images. This includes both the system images, for example asacx-sys-9.1.1-1.pkg, and also the boot images, for example asacx-5500x-boot-9.1.1-1.img. Using Image Manager, you cannot add any CX images to the Image Manager repository and cannot push any CX images to the device .

Handling of SSL VPN Images

Image Manager only reliably copies SSL VPN images to the ASA device. No configuration or activation commands are added for SSL VPN images by Image Manager. The configuration of the images must still be done using Configuration Manager.

The following files are not managed in Image Manager and have to be configured and deployed from Configuration Manager as in earlier versions of Security Manager:

- CSD Configuration XML
- AnyConnect Client Profile files
- DAP Configuration XML
- Full Customization XML files

After the SSL VPN images have been copied to the device using Image Manager, the remote access VPN policies must be configured in Configuration Manager to make use of these images. The Remote Access VPN policies that must be configured are located at the following paths in Configuration Manager:

- **CSD Package**—Remote Access VPN > Dynamic Access > Cisco Secure Desktop group box
- **HostScan Package**—Remote Access VPN > Dynamic Access > Cisco Secure Desktop group box
- **Anyconnect Image**—Remote Access VPN > SSL VPN > Other Settings > Client Settings tab
- **Plug-ins**—Remote Access VPN > SSL VPN > Other Settings > Plug-in tab

The SSL VPN binary files must be present on the device flash before you reference them in VPN policy. If not, Security Manager will present an activity validation warning informing the user of the preference to use Image Manager to push these files reliably to the device before deploying the configuration. If the user ignores the activation warning and goes ahead, Configuration Manager defaults to the old behavior and pushes the images or files as was done in the earlier versions of Security Manager before deploying the configuration referring to these files. But the user cannot leverage the following advantages of using Image Manager for copying these files:

1. Capability to use external disks like disk1 to copy the files. Configuration Manager only copies the files to disk0 and does not recognize or support external disks.
2. Image Manager preempts errors during the image copy by validating that there is enough free space on the disk to copy the selected images and does not allow creation of a job unless there is sufficient space is to copy the images. User can make space by using the Image Manager to delete unwanted images.

Note Image Manager does not validate the compatibility of the SSL VPN files that are pushed to the ASA. But Configuration Manager will complain when incompatible files are referenced in the Remote Access VPN policies.

Administrative Settings for Image Manager

Image Manager introduces new administrative settings. These administrative settings must be configured as part of Configuration Manager.

Configuring Cisco.com Certificates

Beginning with version 4.4, Security Manager has a certificate trust management feature. This feature helps you with improved handling of Cisco.com certificates. For detailed documentation of this feature, refer to Certificate Trust Management, page 10-18.

To configure administrative settings for Image Manager, do the following:

Step 1 Go to **Configuration Manager > Tools > Security Manager Administration**.

The Cisco Security Manager - Administration page appears.

Step 2 Configure workflow settings:

Tip Refer to the Configuration Manager documentation for workflow control setting information.

a. Select **Workflow**.

b. To require that Install Jobs be approved explicitly by an assigned approver before they are installed on devices, select **Require Deployment & Install Image Approval**. If you select this option, make sure you configure appropriate email notification settings. For more information, see the Workflow Page, page 11-75.

Note To allow the submitter to approve deployment jobs, select **Submitter can Approve Deployment Job**.

c. Click **Save.**

Step 3 Configure debug settings:

a. Select **Debug Options** and from the drop-down list for Image Manager Debug Level, select the debug level you want.

Tip The levels include: Severe, Error, Warning, Info and Debug. The default log level is Error.

Note The log files are stored as follows:

- The server logs are located at: ***%NMSROOT%\MDC\log\operation\vmssharedsvcs.log*** and ***%NMSROOT%\MDC\tomcat\logs\stdout.log***
- The client logs are located at: ***<Client Install Dir>\logs*.log***

b. Click **Save.**

Step 4 Configure Cisco.com credentials:

a. Select **Image Manager**.

The Image Manager page appears.

b. If you have credentials to connect to Cisco.com configured already at Tools > Security Manager Administration > IPS Updates > Update Server and wish to reuse the same for Image Manager, then check the **Use IPS Updates Settings** check box. This is also the default behavior.

Note Only Cisco.com is supported, not Local Server.

c. On the Image Manager page, deselect the **Use IPS Updates Settings** check box if you want to specify a set of credentials for Image Manager explicitly.

The fields on the Image Manager page become operable.

d. Complete the following fields:

- Username
- Password
- Confirmation

e. Optionally, complete the Proxy Server Settings to configure a proxy, if required.

1. Select the **Enable Proxy** check box.
2. Complete the following fields to define the proxy:
   - IP or Hostname
   - Port
   - Username
   - Password
   - Confirmation (of Password)

f. Click **Test Connection** test connectivity to Cisco.com with the configured settings.

g. Click **Save**.

Step 5 Configure Purge Interval for Image Install Jobs

a. Select Image Manager.

b. Enter a purge value to specify how many days should pass between purges, in the Purge Jobs Older Than field.

Note Pressing the Purge Now button immediately purges the Image Installation jobs satisfying the Purge Interval criteria.

Step 6 Configure Image Backup Settings

a. Select Image Manager.

b. To include the repository as part of the standard backup, select **Include Repository**.

Caution Ensure that you have sufficient hard disk space on the Security Manager server as the image files consume a lot of space.

Note You can click the Reset button to reset the values to the last saved values before the current change.

c. Click **Save**.

Step 7 Click **Close** to close the Administration window.

Bootstrapping Devices for Image Manager

The bootstrapping in Image Manager is essentially the same as that which you perform in Configuration Manager for ASA devices.

To bootstrap a device for Image Management, do the following:

Step 1 Configure HTTPS on the device(s) to manage ASA in Security Manager.

a. Ensure that the HTTP server is enabled.

b. Add the Security Manager server IP address as an allowed host for HTTP management on the device.

Step 2 Ensure that the configuration register setting is set to boot with the image list in the running configuration.

a. Register value: 0x1,0x3,0x5, 0x7, 0x9

Note Register value: 0x1 is the recommended setting.

b. Do not set to boot to **rommon** mode. (Otherwise device will not be rebooted and the image upgrade will be aborted.)

Step 3 In Security Manager, go to Tools > Security Manager Administration > Device Communication > SSL Certificate Parameters. In the SSL Certificate Parameters area, set PIX/ASA/FWSM Device Authentication Certificates to Do not use certificate authentication.

Step 4 Ensure that there is sufficient space in the flash memory of the device(s) to hold the images you intend to load.

Tip If necessary, you can delete other images you do not intend to use from the device(s).

Step 5 We recommend that you unmanage the Boot-Image/Configuration policy for ASA, as follows:

a. In Security Manager, navigate to Tools > Administration > Policy Management.

b. Uncheck the Boot Image/Configuration policy selection.

Note Image Manager configures boot image and ASDM image as part of the image installation job. So, if the Boot Image/Configuration policy is not unmanaged, then any configuration deployment after the image installation will remove these boot commands added by Image Manager. To prevent this, the Boot Image/Configuration policy should be unmanaged in Security Manager. This can be done from Security Manager administration settings -> Device Exception Settings -> Firewall Policies node.

Step 6 We recommend that the device not be set as a priority monitored device in HPM.

Step 7 Ensure that all configuration changes on the device are submitted and deployed.

Working with Images

Image Manager provides access to images on Cisco.com as well as images on your network. When an image shows a location of Repository, it means that image has already been downloaded (either from Cisco.com or from a local file system). Conversely, an image that shows location as Cisco.com has not been downloaded into the repository. Navigating to the Repository Images in the Images section of the selector enables you to examine a list of all the images. You can also filter, sort, and search the images available. Filtering, in particular, is a good way to navigate within Image Manager. Beginning with all images, you can use the headings in the main repository view to locate images by a variety of attributes including name, version, and type.

Image Manager does not manage ASA-CX images. Any CX images available on Cisco.com will not be shown in Image Manager for download. You also cannot add any CX images from the file-system.

Note Only images that are downloaded to the Image Repository can be used for image upgrade jobs.

Note Beginning with Security Manager release 4.4, when Security Manager contacts Cisco.com to update images or to check on the availability of image updates, an additional certificate validation is performed. The update or download fails if you have not accepted the most recent certificate. You must retrieve,

view, and accept the most recent certificate before you can proceed with other operations. For more information on the certificates, see Managing Device Communication Settings and Certificates, page 9-4.

This section contains the following topics:

- View All Images, page 73-10
- Download Images to the Repository, page 73-11

View All Images

When you first open Image Manager, or when you select All Images from the selector, the system displays a complete list of images. This list includes both those images in the repository as well as those on Cisco.com (which have not yet been downloaded). Some of the VPN image files are bundled with the Security Manager installation and are shown in the repository from the first time onwards. Image Manager will display a warning about credentials not being configured when the Image Manager client is launched for the first time, or until the credentials are configured under Security Manager administration settings for Image Manager.

Note In earlier releases of Security manager, only the prepackaged SSL-VPN images already existing in the Image Manager repository could be seen. If you do not have a repository connection, beginning with Security Manager release 4.4, on a freshly installed Security Manager, Image Manager shows not only the prepackaged SSL-VPN images in the repository, but also lists supported ASA images available on Cisco.com. The prepackaged files are available at: `CSMRoot>\MDC\athena\ccometadata`. Thus, even if you do not have initial connectivity to Cisco.com, you can view the latest images that are available at the time of release of the Security Manager. You have to have Cisco.com connectivity and should configure credentials to Cisco.com to either check for the latest updates on Cisco.com or to download the images from Cisco.com, or both. This prepackaged information about image availability enables users not having Cisco.com connectivity to still view the latest images available on Cisco.com (at least the ones published on Cisco.com by the Security Manager release). This also lets you view compatible images for a particular device type/platform.

This view can be re-ordered by any of the listed image attributes. For example, you can list the images by size. The attributes that you can sort on include:

- Download State - This is the first column and is shown as icons. The icons are actionable, and you can double-click the icon in this column to start the download of an image from Cisco.com, or to abort an ongoing image download, or to delete an image from the repository. Note that the icons change during each of these actions. (A green arrow indicates an image on Cisco.com, a red cross indicates an image that has already been downloaded, and another icon indicates that a download is in progress.)
- Image (name)
- Type
- Version
- Location
- Size
- Description
- Comments (you can add and edit comments for an image).

To view all images, do the following:

Step 1 Check for new images available on Cisco.com

a. Configure the credentials for reaching Cisco.com by navigating to Tools > Security Manager Administration > Image Manager.

b. In the upper right corner, click the double arrow **Check for Updates** icon.

c. Ensure that the CCO account has permissions to download crypto images. Otherwise, navigate to the link and accept the agreement, and then retry the operation.

The system displays "***Updating***" while it checks for updates. When finished, it reads: ***Last updated at: <timestamp>***, and you can view the new images available in the All Images view.

Step 2 If you have not already accepted the most recently issued Cisco.com certification, the system notifies you that you must retieve, view, and accept the latest certificate before any communication with Cisco.com by Image Manager can occur.

Step 3 Click **All Images** in the selector.

The system displays the image list.

Step 4 To re-order the list, click on any of the column headings.

The list of images is reordered according to the selected attribute.

Step 5 To filter the list, use the Image Manager's search window to enter a key string. For example, you could enter the digits of a version number.

Note Also, you can use the filter settings in some of the column headings to filter the list shown.

Download Images to the Repository

You can download images to the repository either from Cisco.com or from a local file system.

Note Beginning with version 4.4, Security Manager has a certificate trust management feature. This feature helps you with improved handling of Cisco.com certificates. For detailed documentation of this feature, refer to Certificate Trust Management, page 10-18. You must have accepted the latest certificate from the image download site on Cisco.com to proceed. The certificate of the site from which the image is to be downloaded may be different from the site that is contacted for "Check for updates" to obtain the latest meta-data information about images. Thus, even if you have accepted the certificate from the "Image Meta-data Locator" URL, the image download may fail with an error to accept the certificate of the image download URL. You must retrieve and accept the certificate from the download URL given in the error message to proceed with the image download.

Note Beginning with version 4.9, Security Manager mandates you to read and accept the End User License Agreement (EULA) before you can proceed to downloading images from cisco.com.

Tip Images can also be downloaded from the Compatible Images tab. For details see Manage Images on a Device, page 73-17.

To add an image file to your Security Manager repository, do the following:

Step 1 To download images from Cisco.com, do the following:

a. From the All Images view double-click the **Start Download** icon in the first column.

Tip Ensure that the credentials to Cisco.com are configured and you have authorization to download the images.

Note Image Manager displays an error message if the image to be downloaded already exists in the repository. The system skips the download when the file name and checksum are identical.

The Downloads window appears showing the progress of the download.

Tip The progress icons may change as the download progresses. A green check icon with the word Deployed indicates success. A red X icon indicates failure. In case of failures, you can view the cause of the failure for that image in the Downloads window. You can double click the message to view the complete details of the error.

b. When completed, select Repository Images and view the image in the listing.

Tip You can also select multiple images and download them all at once by right-clicking them and using the context-sensitive menu.

Tip By sorting the list on Update Time, you can easily view the most recent image.

Step 2 To download images from a local file system, do the following:

a. From the toolbar in the Repository Images view, click the **Download image from file system** icon (found on the far left).

The Download from File System dialog box appears.

b. Use the Browse feature to select the Import location and to select the image to be imported.

c. Click **OK**.

d. Click **OK** in the Download image from file system dialog box.

e. Observe the progress of the download.

Note If the image to be downloaded already exists in the repository, the system displays an error.

f. When completed, select a device group in Security Manager and view the image in the listing.

Tip By sorting the list on Update Time, you can easily view the most recent image.

g. Alternatively, you can download an image file using the drag-and-drop method. For example, you can drag one or more files from your desktop and simply drop it on the Image Manager application.

Working with Bundles

Bundles are groups of compatible images that you define. You can use bundles to simplify repetitive operations, by grouping images that are pre-validated to work together as a logical group. For example, you might define bundles that reflect ASA and ASDM pairs to ensure that you deploy both types in a single operation. The following types of images can be part of a bundle:

- ASA system software
- ASDM image
- VPN images (including csd, AnyConnect, Hostscan)
- Plug-ins (including rdp, ssh, ica, owa, and others)

Multiple system software images cannot be included in the same bundle.

Caution The Image Manager does not stop you from adding incompatible images as part of a bundle. You must determine this compatibility. The ASA and ASDM Compatibility matrix is located at: http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html#wp42231.

Tip Throughout Image Manager there are operations where you can choose to apply an image, multiple images, or a (predefined) bundle of images.

This section contains the following topics:

- Creating Bundles, page 73-13
- View Images by Bundle, page 73-14
- Renaming Bundles, page 73-15
- Deleting Bundles, page 73-15
- Deleting Images from Bundles, page 73-15

Creating Bundles

You can define bundles of images to simplify your Image Manager. Bundles are particularly useful when you have a group of images upon which you regularly operate.

To create a bundle, do the following:

Step 1 From the Bundles heading in the selector, click the **Add Bundle** (plus sign) icon.

Step 2 In the Create Bundle dialog box that opens, enter the name for the new bundle.

Step 3 Click **OK**.

The bundle is listed under the Bundles heading in the selector.

Step 4 From the Images section of the selector, select an image to be bundled. Then click on the Release Notes tab. Finally, examine the compatibility table in the applicable release notes to ensure there are no conflicts with the other images to be bundled.

Caution The Image Manager does not stop you from adding incompatible images as part of a bundle. It is the responsibility of the user to determine this compatibility. The ASA and ASDM Compatibility matrix is located at: http://www.cisco.com/en/US/docs/security/asa/compatibility/asamatrx.html#wp42231.

Step 5 With compatibility determined, drag and drop each image onto the bundle.

Note Multiple system software images cannot be included in the same bundle.

Tip You can select a range of images to drag and drop by selecting the first image in the range and then, with the Shift key pressed, selecting the last image in the range. You can select multiple images by clicking those images while keeping the Ctrl key pressed. You can also select a range of images and then add additional images to your selection by using the Ctrl key method. To move multiple images to a bundle, drag using the right mouse button.

View Images by Bundle

You can view the images that have been added to a bundle.

To view the images in a bundle, do one of the following:

Step 1 To view the images contained within all bundles:

a. Under the Bundles heading in the selector, click the top-level Bundles folder.

   All the bundles are listed, together with the images contained within each.

b. You can expand or collapse any of the bundles to make viewing easier. To expand all bundles or collapse all bundles, use the Expand All and Collapse All buttons at the top of the main window.

Step 2 To view the images for a specific bundle:

a. Under the Bundles heading in the selector, select a bundle.

   The list of images for the selected bundle is displayed in the main window.

Step 3 To view a particular bundle:

a. In the Bundles section of the selector, click **Search** (the magnifier icon).

b. Enter the bundle name in the search field under the Bundles banner.

The list of bundles displays only the specified bundle.

Renaming Bundles

Bundles can easily be renamed to provide better organization or to more accurately reflect the contents of the bundle.

To rename a bundle, do the following:

Step 1 From the Bundles heading in the selector, select a bundle.

Step 2 Right-click the bundle name and, from the drop-down list, select **Rename Bundle**.

The Rename Bundle dialog box appears.

Step 3 Type the new bundle name.

Step 4 Click **OK.**

The new bundle name appears in the selector under Bundles.

Deleting Bundles

Bundles that are no longer needed can be deleted.

To delete a bundle, do the following:

Step 1 Select the bundle.

Step 2 From the Bundles heading in the selector, click **Delete** (the red X icon). Alternatively, right click on the selected bundle and select **Delete Bundle**.

Deleting Images from Bundles

If you wish to change the contents of a bundle, you can delete any of the images that are defined as part of a bundle.

To delete images from bundle, do the following:

Step 1 Select the bundle.

Step 2 Right-click on the image(s) to be removed.

Step 3 Select **Delete Image from Bundle**. Alternatively, click the Delete button at the top of the table.

Working with Devices

The following topics explain how to work with devices in Image Manager.

Note For a cluster, only the control unit node supports the download of files from storage.

This section contains the following topics:

- Viewing Device Inventory, page 73-16
- Manage Images on a Device, page 73-17
- View Device Memory, page 73-18
- Configuring the Image Install Location, page 73-19

Viewing Device Inventory

You can use the Device Summary page to quickly view the devices on your network and their attributes.

Within the selector panel on the left is an area called Devices. From that area you can display all devices by selecting All (or you could select a location or group of devices you have defined). After you have chosen the scope of device selection, the corresponding devices are displayed in the upper panel of the Device Summary page. The upper window of the Device Summary page displays the following attributes, as applicable, for each device:

- Device Display name
- Mode (for example, Standalone, Active-Active. Active-Standby, Cluster)
- System SW Version
- ASDM Version
- AnyConnect Version
- Secure Desktop Version
- Hostscan Version

The Device Summary table includes a Mode column. This column specifies such modes as Cluster, Standalone, Active-Active, and Active-Standby.

For configurations in which multiple physical devices are grouped together, as in failover and cluster configuration, each physical unit/ member has its own file system. And these file systems can be different. The details of the file systems of each physical device/ member are viewable within Image Manager.

Details of individual cluster members, including storage and image status, are shown in the Security Manager user interface. During discovery of image management inventory data, details are discovered regarding each cluster member's storage and running image details.

When a Failover or Cluster device is selected in the Device Summary page, the individual physical members in the group are displayed in the middle Device View table. The Device View table for Cluster device displays the following information about the cluster members:

- **Name**—Device or cluster member name.
- **ID**—Cluster Member ID.
- **Status**—Role of the member in the cluster. For example, Cluster Master or Cluster Slave.

- **Serial Number**—Serial number of the cluster device.
- **Running OS Version**—Version of OS on the particular member.
- **CCL IP**—Cluster Link IP address.
- **CCL MAC**—Cluster Link MAC address.
- **Site ID**—Site ID of the Cluster device.

The Device View table for a *Failover device* has columns that include Name, Status (for example Standby or Active), Serial Number, RAM size, and Running OS Version. The Failover Device table lists these elements by Primary and Secondary devices for the failover pair nodes.

When you select a particular device in the device summary page, then the lower window displays the following tabbed pages for the details of that device.

- **Summary**—Display Name, Device Type, IP Address, Hostname, Domain Name, Serial Number, Running OS Version, Target OS Version, RAM, Failover Mode, Image Install Location
- **Compatible Images**—Images compatible to the device—Image, Type, Version, Location, Size, Description, Comment.
- **History**—Chronological view of Image Installation Jobs and Configuration Deployment Jobs that have been executed on the device—Job Name, Changed By, State, Last Action, Tickets

When you select a particular member of a failover or cluster device in the middle Device View, the lower window displays the following tabbed details for that physical device.

- **Summary**—Running OS Version, Target OS Version, RAM
- **Storage**—The number and capacity of flash memory units. Name, Size, Path, Type, Disk Usage
- **Running Images**—The images presently operating. Name, Type, Version, Path, Size

Manage Images on a Device

You can use the Image Management tool to review, download, and remove the images on the ASA device(s) you select.

To review, download, or remove ASA images on a device, do the following:

Step 1 Select a device group from the Devices area of the selector panel.

The main window displays the Device Summary. The Device summary lists the devices and the associated system software versions.

Tip Alternatively, you can select the search function (magnifier icon) from the Devices banner and then enter the device name in the search field that appears.

Step 2 From the upper pane of the Device Summary page, select a device.

Note If a particular device is part of a cluster, you can navigate through the cluster to view device details.

The lower pane displays details of the selected device.

Step 3 Select the Storage tab in the lower pane and then check the amount of free space listed under Disk Usage.

Note If a particular device is part of a cluster, you can navigate through the cluster to view device storage details.

Tip The device may have more than one storage area, for example, disk1. Be sure to scroll down to see secondary (flash) storage capacity.

Step 4 Note the available disk space on the device.

Step 5 To remove one or more images from the device to free up space, select one or more images in the Storage Tab and click **Delete** at the top of the Storage Tab.

Tip Alternatively, you can select one or more images, right-click and click Delete.

Tip If you delete an image that is currently active and is being referenced, Image Manager displays a warning message.

Step 6 To download an image from the device, select the image and click **Download** at the top of the Storage tab. Select the location on the local file system to which to download the image and click **OK**.

Note For a cluster device, download of images is only supported on the Cluster master. Similarly, for a failover device, download of images is only supported on the active device in the pair.

A dialog appears showing the progress of the download from the device. After the download is completed, the downloaded image is shown in Explorer.

Step 7 Select the Compatible Images tab in the lower pane.

The system displays images that are compatible to the device.

Step 8 To install the compatible image(s) onto the device do the following:

a. Select the image to be added to the device.

b. Double-click the download icon.

The image is downloaded to the repository.

c. Select the image and then select **Install** from the context menu.

The Install wizard appears and the system installs the image. Please see Install Compatible Images on Devices, page 73-30 for details.

View Device Memory

You can use Image Manager to determine the memory capacity and application of a device in your network.

Note Only physical devices can display memory capacity, clusters do not.

To view the details of the memory on a device, do the following:

Step 1 From the Devices area of the selector panel, choose the device to examine.

Details of the selected device appear in upper window of the Device Summary page.

Step 2 In the upper panel, examine the RAM listing.

Caution Image Manager warns you if there is insufficient RAM on a device to load the new image. However, the system does not stop you from performing such an image upgrade. (This is in contrast with Configuration deployment, where the deployment job stops if there is insufficient RAM.)

Configuring the Image Install Location

ASA devices have a default flash (disk0) where all the images reside. By default, Image Manager copies images to disk0 of the ASA device. When the ASA device is configured with an external disk (that is, disk1), Image Manager allows you to choose between the two disks, disk0 or disk1, when loading images on the ASA device.

Note The capability to load images on an external disk is very useful for storing large images such as those for AnyConnect and CSD, as disk0 can run out of space quickly with even a few of these larger images.

For configuring Image Manager to use the external disk, do the following:

Step 1 Select the device from the Devices area of the selector.

Step 2 View the Summary information on the right pane.

The available disks on the device are listed in the Image Install Location drop-down list. For a device with external disk, disk0 and disk1 would be listed.

Step 3 Select the external disk, disk1, from the Image Install Location drop-down list and click **Apply**. For a user context device, you can select shared or private label to apply the default Install location.

All future image installation jobs for the device will load the images to disk1.

Tip The configuration of the external disk can be verified by performing an image install operation. After the job is complete, view the contents of disk1 in the Storage tab of the device in Image Manager. It should list the newly installed images.

Note For cluster and failover devices, and multi-context devices, if all the physical member devices do not have the disk that is selected as the Image Install location, then there will be a validation error when you try to copy or install images. You need to select the Image Install location to be a disk that is present on every member device to proceed with copying or installing images.

About Image Updates on Devices Using Image Manager

How does Image Manager update images on an ASA device?

Image Manager follows the standard documented procedure to upgrade the stand-alone ASA devices with several built-in checks to ensure reliable image upgrade. Please refer to: http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a008067e9f9.shtml#maintask2 for the procedure for image upgrade.

Note You must have accepted the latest Cisco.com certificate to enable Image Manager to interface with cisco.com. You must accept the certificate from both the "Image Meta-data locator" site and the download site of the images to start downloading images successfully (see Image Manager Page, page 11-41).

Image Manager uses the HTTPS protocol to copy images to the ASA device, performs configuration changes to activate the new image (ensuring fallback to the older image in case of any error), and finally reloads the device if required, with the new image.

How does Image Manager update images on an ASA configured for failover?

Updating the images in an Active/Standby failover pair is accomplished by creating an image upgrade job on the active device of the pair, and then running the image upgrade job.

Image update on an Active/Active failover pair is not supported in Image Manager. The Active/Active failover pair has to be converted to Active/Standby by making all the failover groups active on one unit, and the corresponding failover groups standby on the other unit. Only then can Image Manager update the image on the pair of devices.

To upgrade devices in an Active/Active failover pair:

1. Manually convert the pair to active/standby by forcing all the failover groups on one device to be **active** and on the other device to be **standby**.

   Note Do not discover the devices in Security Manager.

   Note For additional details on how to convert an active/active failover pair to active/standby, see http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a0080b20f35.shtml#Actact.

2. Create an image upgrade job on the active device of the pair, and run the image upgrade job.
3. After the upgrade has occurred, manually convert the pair back to active/active configuration, as existed before the upgrade, by making the required failover groups active on one unit and the remaining failover groups to be active on the other physical unit.

4. Rediscover in Security Manager only the device inventory for the unit that was converted to standby.

Image Manager follows the upgrade procedure as detailed at: http://www.cisco.com/en/US/products/ps6120/products_configuration_example09186a0080b20f35.shtml. The image is copied to both the units and then configuration change is done to activate the image that is synced to both units. First the standby is reloaded via the active unit and after ensuring that the standby has been upgraded successfully to the new version, the current active is reloaded. After both the units are upgraded to the new version, the failover pair or cluster upgrade is marked successful.

Note During the current active reload and until the standby ASA takes over, the traffic going through the failover pair will be impacted.

There are restrictions for image upgrade in failover ASA pair. We recommend that, while performing image upgrades on a failover ASA pair or cluster using Image Manager, you ensure the following restrictions are satisfied:

- The two units in a failover configuration should have the same major (first number) and minor (second number) software version.
- **Maintenance Release**: You can upgrade from any maintenance release to any other maintenance release within a minor release. For example, you can upgrade from 7.0(1) to 7.0(4) without first installing the maintenance releases in between.
- **Minor Release**: You can upgrade from a minor release to the next minor release. **You cannot skip a minor release.** For example, you can upgrade from 7.0 to 7.1. Upgrading from 7.0 directly to 7.2 is not supported for zero-downtime upgrades; you must first upgrade to 7.1
- **Major Release**: You can upgrade from the last minor release of the previous version to the next major release. For example, you can upgrade from 7.9 to 8.0, assuming that 7.9 is the last minor version in the 7.x release.

How does Image Manager update images on an ASA cluster?

Image updates follow the procedure previously established for hitless upgrades that ensures that all members of a cluster are upgraded to a new version in a single user operation without affecting traffic flow. During an image upgrade:

- Slave cluster members are first loaded with the new image from the master. Images are copied to all members of a cluster while being connected only to the master. Such propagation through a cluster does not require switchovers of each device to master status and, thereby, minimizes traffic disruption.
- Configuration is changed on the control unit to add the boot command to load with the new image. The configuration, once changed on the control unit, automatically gets synced on all the data units.
- All the data units reboot with the new image sequentially via the control unit.
- All the data units come online and rejoin the cluster.
- The control unit is then made into a data unit (with the next data unit taking over the control unit role).
- Via the new control unit, the old control unit is reloaded with the new image.

This procedure for image update followed by Image Manager ensures minimal switchovers and minimal disruption of traffic.

Device State Changes During and After Image Update

Image upgrade is a critical operation and hence there is a need to depict visually, and inform users of, all image update operations. Thus, three new device states have been introduced:

- **Upgrade In Progress**—The device is put into this state whenever an image install job starts on the device. This state is automatically reset by the system after the image update operation is completed on the device.
- **Maintenance**—The device is put into Maintenance state when the image install job fails on the device and the device becomes unreachable after the image install operation. You need to manually reset this state to normal/operational state after taking necessary steps to bring back the device online by manually correcting issues due to upgrade or by rolling back the image.
- **Configuration Required**—For certain cases of image upgrade (like from ASA 8.2 to ASA 8.3), there are significant changes in the device configuration as part of the image upgrade which renders the policy configuration model in Security Manager incompatible with the device configuration. In such cases, even if the image upgrade operation may be successful, you must perform some operations, like rediscovery of device, to ensure that after upgrade, Security Manager's configuration policy model and the device configuration are in tandem. Thus, after an image upgrade, if some additional configuration is required to be made in Configuration Manager to make the device operational, the device is put into *Configuration Required* state. Even in the case where VPN images are deployed using Image Manager, the device is put into *Configuration Required* state since this requires the user to configure these images in the VPN policies using Configuration Manager. The *Configuration Required* state indicates that there are changes to be done in Configuration Manager to make the device functional in Security Manager after the image update operation. You can make the suggested changes and, after you are satisfied with your configuration changes, you can manually bring the device back to *Operational* state.

Note Refer to the Troubleshooting Image Management, page 73-37 for more scenarios that can put the device in the Configuration Required mode.

Whenever the device state is changed to any of these three states, the state is indicated in the device selector with an explicit icon. This change in device state can be seen in both Configuration Manager and Image Manager. The normal state of the device when there is no image update operation on the device is the *Operational* state.

Tip To manually reset the device state to normal or *Operational* status, select the device(s) in the device selector in Configuration Manager or Image Manager, right click, and select Make Device Operational.

This section contains the following topics:

- Validating a Proposed Image Update on a Device, page 73-23
- Using the Image Installation Wizard to Install Images on Devices, page 73-26
- Install Bundled Images on Devices, page 73-30
- Install Compatible Images on Devices, page 73-30
- Install Images on Selected Devices, page 73-31

Validating a Proposed Image Update on a Device

You can validate the image update job on one or more devices prior to actually performing it. The following list details the various validations that are performed:

- Insufficient disk space on the ASA device to accommodate the selected images.

  An error is displayed in this case. You must navigate to the Storage tab for that device and delete one or more images to make space. Then retry the upgrade validation operation.

  Note The disk space on each of the members in the Cluster and both active and standby units in a Failover are evaluated for sufficient space to accommodate the selected images. When a single member or device does not have sufficient space, an error is displayed and you cannot proceed to creating a job on the device. You must navigate to the Storage tab for that particular member and make space by deleting one or more unwanted images.

- Insufficient RAM on the device to run the new image as recommended in: Release Notes for the Cisco ASA 5500 Series, 8.4(x).

  Note Image Manager will warn you if there is insufficient RAM on the device to load the new image. However, this will not stop you from performing an image upgrade. This is in contrast with Configuration deployment, wherein the deployment job stops if there is insufficient RAM.

- If the flash device (disk0 or disk1) that is selected for Image Install Location is not present on any of the devices/members in cluster/failover setup, then and error is displayed and the job is aborted.
- Configuration changes that have been submitted but not yet deployed to the device. These changes need to be deployed before starting the image update job. Otherwise, the configuration changes may become incompatible with the upgraded image version on the device.
- Warn if the selected image(s) is (are) incompatible with the device type, for example, if non-SMP images are selected for ASA 5585 device types.

  Note This warning only occurs when you are using the drag-and-drop method. For other flows, the incompatible images/devices are filtered out in step 2 of the Image Install Wizard.

  Note This validation was skipped in Cisco Security Manager 4.3 if Check for updates is not performed due to the unavailability of meta-data information about images on Cisco.com compatible with MDF IDs. In Cisco Security Manager 4.4, the meta-data information are prepackaged with Cisco Security Manager install and hence even if Check for updates is not performed, Image manager will validate the compatibility of images for device types and warn the user when an incompatible image-device combination is chosen.

- Warn if the device is being updated to a version that is unsupported on Security Manager.
- Warn if the new image version is the same as or lower than the version running on the device.
- An image upgrade to ASA Version 8.3 from any lower version would require the device to be re-discovered in Security Manager. There were major changes in the NAT configuration introduced in ASA version 8.3 that are incompatible with previous ASA versions. Likewise, there were major

changes to the NAT policy model in Security Manager for Version 8.3. Hence, when a device is upgraded to ASA 8.3, that device is put into the Configuration Required state to indicate to the Configuration user that some changes are required in Configuration Manager to make the device operational. After rediscovering the device in Security Manager, the user can right-click the device in the device tree and select Make Device Operational to bring it back to normal state.

- An image upgrade from ASA Version 8.3.x to 8.4.2 or higher also requires rediscovery of the device in Security Manager because of incompatible changes in PAT configuration in ASA Version 8.4.2. In this case also, the device state is changed to the Configuration Required state after the image upgrade.
- An image upgrade from ASA Version 8.x to 9.0.1 or higher requires rediscovery of the device in Security Manager because of incompatible changes in Unified access rules, inspection, and NAT rules.
- An image upgrade from ASA Version 7.x to 8.x introduces major changes in the SSLVPN configuration both on the device and in Security Manager. Because of these incompatible changes, the device must be deleted from Security Manager after the image upgrade and then added again. The device is put into the Configuration Required state to notify the user to act upon these warnings.
- Any VPN image, such as an AnyConnect, CSD, or Hostscan image, being loaded using Image Manager warns the user if the existing VPN image was part of a shared policy being assigned to the current device or other devices. A warning is issued to also copy the new image to all the devices to which the shared policy is assigned so that the shared policy can be updated for all the devices seamlessly without the loss of policy sharing.
- Warn when the standby unit in a failover pair in not reachable. This is an error that will cause the job to be aborted.
- Warning for upgrade of Active/Standby failover pair. The version being upgraded should comply with the recommendations listed at: ASA/PIX: How to Use the CLI to Upgrade the Software Image on a Failover Pair.

  Note The same warnings are also applicable for an ASA in cluster configuration.

- Warning for upgrade of Active/Active failover pair that the image update job will be aborted unless the pair is converted to Active/Standby (that is, if all the failover groups are active on one physical unit).

Note Additional validations are required for a "device" that is a designated control unit. In addition to a check that is similar to that for Active/Standby Failover, there is also a check that the cluster image is compatible with supported platforms. A cluster cannot be downgraded to a version less than 9.x.

To validate an image installation, do the following:

Step 1 From the File menu, select **Validate.**

The Validate Image Assignments window opens.

Step 2 Click **Add Assignment**.

The Image Assignment window opens.

Step 3 In the Image Assignment window, from the drop-down list, select either:

- Select Images and Assign to Devices

- Select Devices and Assign to Images

Tip The same result is obtained by assigning images to devices or by assigning devices to images.

Step 4 Select one or more items (images or devices) by moving them to the window on the right.

Tip You can use pre-defined bundles rather than images by clicking **Bundles** and selecting the bundle.

Step 5 Click **Next** and assign the other item (images or devices).

The Confirm Assignments window appears.

Step 6 Examine and confirm the assignments that you have specified.

Tip You can continue to add or remove assignments, as required.

Step 7 Click **Finish.**

The Validate Assignments window appears.

Step 8 Click **Start Validation**.

The validation status of the assignment is shown in the Validation column. It shows either Warning or Successful or Error.

Step 9 Click on the display of Warning or Successful or Error.

A lower pane window opens.

Step 10 According to the validation status, do one of the following:

- **Error**—Examine the potential reasons for the error and correct as necessary.
- **Warning**—Examine the potential reasons for the warning and correct as necessary.
- **Successful**—No action required.

Tip Be sure to display and examine all the errors or warnings by using the window slider bar on the right.

Step 11 After addressing the warnings or errors shown, you can proceed to create the job using the Image Install wizard.

Step 12 Optionally, you can right-click on an assignment element in the Validate Image Assignments window and modify or delete it before proceeding.

Step 13 You can also right-click on an assignment and select **Copy Table**. This copies the assignment details and the validation status and notes. You can then paste the contents to Notepad or another program as a CSV file to be used as reference.

Using the Image Installation Wizard to Install Images on Devices

You can use this feature to create a job to assign and install images on devices. An assignment is simply an association of an image and a device that defines an installation job.

Note If you have the workflow function enabled, you must perform the additional steps described for obtaining authorization before you can accomplish installation.

Note You can choose to operate upon any arbitrary set of devices.

To create a job to install images on devices, do the following:

Step 1 Go to Files > Open Image Installation Wizard.

The Image Installation Wizard dialog box appears.

Tip The Image Installation Wizard can be invoked in several ways. In addition to invoking the Wizard from the menu, as described here, you can invoke it when you do any one of the following things: (1) install images by the drag-and-drop method; (2) right-click on a device or bundle; or (3) select a device, navigate to the Compatible Image tab, select one or more images from the table, right-click, and select the **Install** option.

Step 2 On the lower left, click **Add Assignment**.

The Image Assignments dialog box appears.

Step 3 From the drop-down list on the top, select whether you want to assign images to devices, or devices to images.

Step 4 Move items (devices or images) from the list on the left to the selected items list on the right. Then click **Next**. You can also select bundles, instead of images, by clicking the Bundles tab.

Tip When pairing images and devices to define assignments, you can proceed with images and then devices, or devices and then images. The order of this selection does not matter.

Step 5 Review the assignment definition in the Confirm Assignments dialog box and click **Finish.**

Tip At this point you can choose to add additional assignment pairs, as desired, by clicking **Add Assignment**.

Step 6 When you are finished defining assignments, click **Start Validation.** Wait until the system displays Validation Complete.

Step 7 Examine the status in the Validation column on the Assignments tab of the Wizard dialog box. It shows either Warning or Successful or Error.

Step 8 Determine what steps are necessary according to the status:

- **Error**—Examine the potential reasons for the error and correct as necessary.
- **Warning**—Examine the potential reasons for the warning and correct as necessary.

- **Successful**—No action required.

Step 9 Right-click on the assignment for more options:

- **Move Up/Move Down**—Select these options for a multi-device job when you want to change the order in which the devices will get updated. This feature can be used to order or sequence the devices when the Install Images to Devices job option is set to Sequential.
- **Delete/Delete All**—You choose these options to remove one or all devices from the image upgrade job.
- **Copy Table**—Use this to copy the warning messages to some text editor or spreadsheet program for reference.
- **Test File copy**—Use this option to check whether files can be copied between Security Manager Image repository and ASA device flash using https protocol.

Tip Be sure to display and examine all the errors or warnings by using the windows slider bar on the right.

Step 10 If you want the installation job scheduled for a particular time, select the **Schedule** tab and specify the date and time.

Step 11 To set properties of the installation job, select the **Properties** tab.

a. Edit the **Name**, if desired. (The default is Image install Job—<timestamp>)

b. Edit the **Description**, if desired.

c. Select a **Ticket ID**, if desired.

Tip Starting from release 4.4, Ticket ID field in Image Manager has been decoupled from Config Manager. It is now just a 'tag' and can be any arbitrary string. Ticket ID field is an editable combo box with auto completion that shows tickets created earlier both in Image Manager and Configuration Manager. Also, there is no dependency on Ticketing mode in Configuration Manager for Ticket ID field. Ticket ID is an optional field and can be left blank. Global search in Configuration Manager also supports tickets used in Image Manager and lists the Image Installation jobs with which the ticket is associated.

d. Set the **On Error** option. (Default is Stop Installation, alternative is Continue Operation.)

e. Set the **Backup Current Image** option. (Default is Yes, alternative is No.)

Tip This is only applicable for system software images

f. Set the **Install images to devices in** option. (Default is parallel, alternative is Sequential.)

g. Select one of three operations:

- Install image and reboot device
- Install image but do not reboot device
- Only copy image onto devices

h. If you are using Workflow, you can optionally configure the following approval options:

Tip These are located in the top frame in Job properties for the job.

- **Action**—
- Approve
- Reject
- Deploy
- Submit

If you reject a job, the status is set to Rejected, after which you discard the job. When you discard a job the status is shown as Discarded and all the job's action buttons are disabled.

If you approve a job, the status is set to Approved. Then, you must click Deploy to start the image upgrade job.

i. Beginning with version 4.12, Security Manager provides an option to select the Storage URL for ASA multi-context devices running the software version 9.6(2) or later. You can select either Shared or Private Storage URL for the selected user context. By default, Shared is selected.

Tip You can the check running commentary by selecting the Details tab. and clicking **Show Progress**.

After you deploy a job, the job status is shown as either Deployed or Failed. The History tab in the bottom pane (for the selected job) only is activated in WF mode and displays one of two job action flows:

- Creating/ Edit-In-use/ Submitted/ Rejected/ Discarded
- Creating/Edit-In-use/ Submitted/ Approved/ Deploying/ Deployed (or Failed)
- **Submit the job**—This is checked by default
- **Approver email**—The email address list of approvers
- **Submitter email**—The email address of the person submitting the job

j. You can change job properties by clicking **Edit.**

k. Refer to Viewing Install Jobs, page 73-34 for additional job viewing options.

Step 12 Click **Install**.

The Jobs page appears and the install job is shown with its status as Deploying.

Note If a schedule was selected for the job, the job state is shown as Scheduled. The job will start deploying at the scheduled time and, at that time, the job state changes to Deploying.

Note If workflow is enabled for Image Install Jobs, then the job state is changed to either Submitted or Edit-in-Use. In this mode, the job can be deployed only after it has been approved. Please refer to Image Installation Job Approval Workflow, page 73-36 for details on the job states in the workflow mode.

Tip You can halt the job by clicking **Abort**. Please see Aborting an Image Installation Job, page 73-35 for important information about aborting an installation job.

You can discard a job before the scheduled run time by clicking **Discard**.

Step 13 When the job starts deploying, notice the change in the state of the devices to *Update in progress* state in the device tree in Configuration Manager and Image Manager. A green progress icon appears beside the device in the device tree.

Step 14 View the details of the job and its progress while it is in deploying state. Please see Viewing Install Jobs, page 73-34 for details.

Step 15 Wait for the job to complete.

The job state is changed to *Deployed* if all the devices are successfully updated. The job state is changed to *Failed* if one or more devices failed in the job.

Step 16 Notice the change in state of the device(s) in the device tree after the job is complete.

If the image update is successful and there are no further configuration changes required in Configuration Manager, the device will be moved back to *Operational* state. If there are configuration changes required on the device after the update, then device is moved to *Configuration Required* state. Clicking on the device in the device tree brings up a balloon tip with details of the state and actions to be taken to move the device state back to *Operational*.

If the image update fails on the device and the device is rendered unreachable during the image update operation, the device is put to *Maintenance* state.

Step 17 Verifying the image update:

a. Click on the device in the device tree in Image Manager.

b. Go to the Summary tab to view the updated Running OS version.

c. Go to the Running Images tab to view the new running images after the image update.

d. Select the device in Configuration Manager.

e. Right-click the device and select **Device Properties**.

f. Notice the new image version updated in Running OS Version field.

g. Go to **Configuration Manager > Manage > Configuration Archive**.

h. From the device CLI, enter ***sh ver***

The updated OS version should be displayed.

i. Select the device in the left device tree.

j. View the Configuration Archive versions in the right pane and notice the latest entry with the Archival Source as *Image Manager*.

k. Select the archived entry and click **View.**

l. Compare this entry with the previous archived version to view the configuration changes made by Image Manager during image update. You can view the boot commands being prepended for the new ASA system software image and/or ASDM image command being added for the new ASDM image.

m. Email notification will be sent to configured recipients with the status of the Image Upgrade job if email notification is configured in Image Manager administration settings.

Step 18 If the device is set to *Configuration Required* or *Maintenance* state after the image update operation, follow the steps below to complete the post-image update requirements to make the device functional from Configuration Manager:

a. Click on the device in the device tree in Configuration Manager or Image Manager.

A balloon tip appears showing the device information.

b. View the contents of the balloon tip. Review the reason for the device being set to the *Configuration Required* or *Maintenance* state. Review also the recommended actions to be taken.

c. Perform the recommended actions.

d. Right-click the device in the device tree and select **Make Device(s) Operational**.

The device is moved to *Operational* state and the icon beside the device in the device tree is removed.

Install Bundled Images on Devices

You can use the Image Manager tool to assign and install compatible images that are grouped as bundles. Bundles simplify repetitive operations and can ensure consistent actions are taken on a group of devices.

To selectively install an image bundle on a device or device group, do the following:

Step 1 Drag and drop the bundle onto a device or device group.

The Install images on devices dialog box appears with device and images in the bundle pre-assigned. If the bundle is dropped onto a device group, all the devices in the group are automatically selected and assigned to the images in the bundle.

Step 2 Investigate any assignment validation errors or warnings listed in the Install images on devices dialog box.

Tip You can choose to schedule the job and also modify the default properties for the job. Please see Using the Image Installation Wizard to Install Images on Devices, page 73-26 for details on scheduling a job and configuring job properties.

Step 3 When your warnings are corrected (or you determine them to be insignificant), click **Install.**

Note Alternatively, you can right-click a bundle and select **Install** to launch the Image Installation wizard with the bundle pre-selected. You can then choose the devices and click **Install** to install the bundle on the selected devices.

Install Compatible Images on Devices

You can use Image Manager to install compatible images on devices.

To selectively install one or more compatible images on a device or device group, do the following:

Step 1 Select a device in the Devices area of the selector and navigate to the Compatible Images tab.

Step 2 Select one or more Repository images in the Compatible Images tab.

Step 3 Right-click a selected image and click **Install**.

The Image Installation wizard appears with the selected images pre-assigned or moved to the right pane in the Select Image page.

Step 4 Click **Next**.

The Select Devices page of the wizard is displayed.

Step 5 Select the devices to which you want to install, and then click **Next**.

The Confirm Assignments page of the wizard is displayed.

Step 6 Confirm the devices and images assignments, and then click **Finish**.

The Install images on selected devices dialog box appears with the devices and image(s) assigned.

Step 7 Click **Start Validation** in the upper-right corner of the Assignments tab. Investigate any assignment validation errors or warnings listed in the Install images on devices dialog box.

Beginning with version 4.9, Security Manager provides an enhanced validation procedure for installing images on devices:

- If you have downloaded the images from CCO using the Image Manager, before installing the images to the device, the serial number of the device is verified for the service contract. If the device has a valid service contract, the image will proceed with the installation or upgrade process. If the device does not have a valid service contract, the image will not proceed with the installation or upgrade process.
- If you have copied the images from the local file system to Image Manager, the service contract validation will not be performed for the device and you can proceed to install the image on the device.

Tip You can choose to schedule the job and also modify the default properties for the job. Please see Using the Image Installation Wizard to Install Images on Devices, page 73-26 for details on scheduling a job and configuring job properties.

Step 8 When your warnings are corrected (or you determine them to be insignificant), click **Install**.

The image installation job is created. Please see Using the Image Installation Wizard to Install Images on Devices, page 73-26 for the remaining steps to monitor the job progress and verify the image update.

Note Alternatively, to install one or more images on a device or device group, you can drag multiple images from the Repository view and drop them onto a device or device group. Then, click **Install** to install the selected image(s) on the selected device(s).

Install Images on Selected Devices

You can use Image Manager to upgrade the images on a set of devices that you select.

To install images on a selected set of devices, do the following:

Step 1 Select a device group in the Devices area of the selector.

Step 2 View the listing of devices in the group in the right pane.

Step 3 Select one or more devices from the list.

Tip Use the Shift and Control keys to select multiple devices.

Step 4 Right-click a selected device and click **Install**.

The Image Installation wizard appears with the selected devices pre-assigned or moved to the right pane in the Select Devices page.

Step 5 Click **Next**.

The Select Images page of the wizard is displayed.

Step 6 Select the images you want to install, and then click **Next**.

Tip You can also select a bundle in the Bundles tab.

The Confirm Assignments page of the wizard is displayed.

Step 7 Confirm the devices and images assignments, and then click **Finish**.

The Install Images on selected devices dialog box appears with the devices and image(s) assigned.

Step 8 Click **Start Validation** in the upper-right corner of the Assignments tab. Investigate any assignment validation errors or warnings listed in the Install images on devices dialog box.

Tip You can choose to schedule the job and also modify the default properties for the job. Please see Using the Image Installation Wizard to Install Images on Devices, page 73-26 for details on scheduling a job and configuring job properties.

Step 9 When your warnings are corrected (or you determine them to be insignificant), click **Install**.

The image installation job is created. Please see Using the Image Installation Wizard to Install Images on Devices, page 73-26 for the remaining steps to monitor the job progress and verify the image update.

Working with Jobs

This section details the set of functions that assists in the performance of the image installation jobs. An image installation job may be immediately run, or it may be scheduled to run at a specified time and date. As Image Manager jobs tend to be time-consuming, the job management functions enable you to perform these operations in the background. Image Manager incorporates an optional ticketing system that enables you to easily locate a job or jobs by means of a unique Ticket ID.

You should understand that the details of a particular job are defined and validated before being run.

This section contains the following topics:

- Viewing Image Installation Job Summary, page 73-33
- Viewing Install Jobs, page 73-34
- Aborting an Image Installation Job, page 73-35
- Retry a Failed Image Install Job, page 73-35
- Roll Back a Deployed Job, page 73-35
- Image Installation Job Approval Workflow, page 73-36

Viewing Image Installation Job Summary

You can use the Image Manager tool to monitor image installation and deployment jobs. You can view the history and status of jobs that Image Manager has performed, as well as the summary, details, or history of any particular job.

Note Comprehensive details of job state changes are available in Configuration Manager (see Job States in Non-Workflow Mode, page 8-4 or Job States in Workflow Mode, page 8-6). For an Audit Report navigate to Configuration Manager > Manage > Audit Report.

To view a summary of image installation jobs, do the following:

Step 1 In the selector, under Jobs, click **Install Jobs**.

The main window displays the Jobs list in the upper pane.

Step 2 Examine the details of the Jobs list, which may include:

- **Name**—The name of the job. By default the name includes a time stamp.
- **Last Action**—Date of the last action.
- **Status**—Whether the job has deployed, failed, or is underway.
- **Changed By**—Who started the job.
- **Description**—Job description.
- **Schedule**—Job schedule.
- **Ticket ID(s)**—Tickets are just tags attached to Image Manager job to track changes. These tickets may be tickets created in Configuration Manager, but are not mandated to be so.

Step 3 Optionally, you can find and select a single job about which more information is then displayed in the lower pane. This includes:

- Summary
- Details
- History

Tip While viewing the lower pane Details tab, you can select a device and then view a log of the job in the lower right pane.

Note These are dockable windows. You can customize the default view.

Viewing Install Jobs

You can view the details associated with a particular Image Management job.

To view the details associated with a job, do the following:

Step 1 In the selector, under Jobs, click **Install Jobs**.

Tip The Status column in the Jobs selector indicates whether a job is Submitted, Approved, Deployed, In Progress, or Failed.

The main window displays the Jobs list in the upper pane.

Step 2 Select a job to examine.

Tip To find a particular job, you can sort the Jobs list by any of the column headings, including Name, Last Action (chronology), Status (e.g., Deployed, Failed), Description, and so forth. You can also find a particular job by using the search window to enter a filtering string.

Note The location of the jobs folder is CSM-ROOT\files\vms\jobs directory.

Step 3 Examine job summary information by clicking **Summary** in the lower pane.

The lower pane displays job summary information including Image Management Job Name, Devices to be Deployed, Devices Deployed Successfully, and Devices Deployed with Errors.

Step 4 Examine job summary details by clicking **Details** in the lower pane.

Details of the Devices, New Image, Old Image, and the device status are displayed.

Step 5 Examine the Commentary on the devices in the job by clicking on the vertical **Commentary** tab on the far right. This shows the progress of the image install operation on the device.

Step 6 Examine the Transcript of the devices in the job by clicking on the vertical **Transcript** tab on the far right. This shows the chronology of the commands executed on the device and their responses.

Step 7 Examine job history details by clicking **History** in the lower pane. This shows the history of the transition in the job state.

Note This information is visible only in the workflow mode.

Aborting an Image Installation Job

You can abort an image installation job by clicking **Abort** from the Jobs page. This option is effective only for multi-device jobs:

Note If the job involves a single device, then Abort will have no effect after the job begins and the job will always run to completion.

- If the Sequential option is selected, then all the devices on which the job has not yet started will be aborted.
- If Parallel is selected, then all the devices till that batch will undergo image upgrade. All devices from the next batch onwards will be aborted.

Retry a Failed Image Install Job

If your attempts to deploy an image to one or more devices fails, you can retry the job. However, you should retry the entire job and not attempt to simply continue from a failed step.

To retry a failed job, do the following:

Step 1 To determine that an installation job has failed, go to the Jobs section in the selector and click **Install Jobs**.

The Jobs page appears.

Step 2 Determine the status of the job in question by examining the Status column.

Tip A green check icon with the word Deployed indicates success. A red X icon indicates failure.

Step 3 Investigate possible reasons for job failure.

Step 4 Select the failed image install job from the job list and click **Retry** from the toolbar atop the upper pane.

The Install Images on Devices window appears. You can observe the validation warnings as you would in a normal install job.

Step 5 As required, you can change the images, devices, schedule, or job properties to be used.

Step 6 From within the Install Images on Devices window click **Install.**

Step 7 Determine that the retried attempt is successful by observing the newly created job.

Roll Back a Deployed Job

You can roll back the changes from a deployed image installation job.

To roll back a Deployed job, do the following:

Step 1 From the job list, select the image install job to be rolled back and click **Rollback** from the toolbar atop the upper pane.

The Install Images on Devices window appears. You can observe the validation warnings as you would in a normal install job.

Step 2 As required, you can change the images, devices, schedule, or job properties to be used in the rollback.

Step 3 From within the Install Images on Devices window click **Install**.

Step 4 Determine that the rollback attempt is successful by monitoring the newly created job.

Image Installation Job Approval Workflow

Image update is a critical operation that has the potential to cause downtime for devices and your network. Hence, change control and management for image install operations is crucial. Change management for image installation jobs is done using the Deployment Workflow framework of Configuration Manager. This ensures that all image installation jobs need to be approved before getting executed or deployed.

To use workflow with image installation jobs, do the following:

Step 1 Enable workflow for image installation jobs:

a. In Configuration Manager, select **Tools > Security Manager Administration > Workflow**.

b. If workflow is not already enabled, select **Enable Workflow**.

c. Select **Require Deployment & Install Image Approval**.

d. Configure the email address of the person responsible for approving the image installation job in the Job/Schedule Approver field. For more information, see Workflow Page, page 11-75.

e. Click **Save** and then **Close**.

f. Launch Image Manager and navigate to **Install Jobs**. Notice the new buttons available in the Menu bar for job state transitions in workflow mode: Submit, Approve, Reject, and Deploy.

Step 2 To create and execute an image installation job with workflow enabled:

a. Use any of the previously documented procedures to create an image installation job. See About Image Updates on Devices Using Image Manager, page 73-20.

Note There is an additional option in the Properties tab to Submit the Job. Check this option to automatically submit the job for approval after creation of the job.

b. Once the image installation job is created, note the state of the job in the Image Install Jobs View.

c. Select the job and click **Submit** to submit the job for approval if automatic submit option was not selected while creating the job.

The job Approver (user with Approver role/privileges) receives an email notification to approve the job.

d. The Approver can log in to Security Manager, launch Image Manager, and navigate to the job.

e. The Approver clicks **Approve** to approve the job after reviewing the details of the upgrade, that is, image being upgraded to, job properties, schedule, and so on.

Job state is changed to *Approved*. The creator of the job receives an email that the job has been approved. Now the job can be deployed.

f. If the Approver is not convinced after reviewing the job details, he can choose to reject the job by clicking **Reject**.

   The job state is moved to *Rejected*. The creator of the job receives an email that job has been rejected. A rejected job will not be deployed.

   Note A rejected job will not be deployed. It can be edited and resubmitted for approval or it has to be discarded.

g. Once the job is approved, the job can be deployed by clicking **Deploy**.

   The job state is changed to *Deploying* and image install job execution is started.

h. If the job is rejected or any other changes are required to be made for a job, the job can be edited by clicking **Edit**.

   The Image Assignments page of the wizard showing all the devices and images is displayed. The user can modify the Job Properties, schedule and even delete some device-to-image assignments and submit the job for approval again by clicking **Submit**.

i. If a job has not started executing, then the user can dismiss the job by clicking **Discard**.

   The job is moved to *Discarded* state. A discarded job does not execute and cannot be edited or moved to any other state.

j. If the changed job is acceptable to the approver, he can approve the job this time and the job can be deployed as mentioned above.

k. Once the deploying job completes execution it will be moved to *Deployed* state if the image installation is successful or it will be moved to *Failed* state if the image installation operation fails.

Troubleshooting Image Management

This section addresses steps you can take to troubleshoot Image Management in response to particular symptoms.

Image installation job might show as failed due to the configured reboot time.

For cluster and failover devices, the reboot time between the standby device and primary device is, by default, set to 15 minutes. If the devices have large configurations, the Image installation job might show as failed due to the configured reboot time. The device will be updated with the image after the configurations are complete. However, due to inconsistencies in reboot time, Security Manager will show the job as failed.

To modify the reboot time, do the following.

- Go to ***%NMSROOT%\MDC\tomcat\vms\athena\WEB-INF\classes.***
- Open the **swimng.properties** file.
- In the following code, modify the reloadTime property:

##MAX_RELOAD_WAIT_TIME for the primary or cluster device

#default time will be 15 mins (15*60*1000)

reloadTime = 900000

- Manually restart the Cisco Security Manager Daemon Manager service.

No data for devices in Image Manager, after Security Manager upgrade. Any one of the following operations that first contacts the device will collect the image inventory for the device:

- Rediscover the device choosing to discover only Device Inventory
- Perform a live deployment to device
- Perform an Image Install operation to the device

Image Download from Cisco.com Fails

- Go to Configuration Manager > Tools > Security Manager Administration.
- Select Image Manager.
- Click **Test Connection** to ensure the server is reachable.
- Check ***%NMSROOT%/MDC/athena/config/****-CCOMetaData.xml*** for error in downloading metadata information for particular MDF IDs.

Update or Image Download fails due to certificate mismatch, unavailability, expiration, or other cause.

- Employ the recommended actions cited by the error message. Use the failed download URL from the error message to retrieve the certificate.
- View the stored certificate at ***%NMSROOT%/MDC/certificates/*.ser*** (Serialized object--file contents are unintelligible and cannot be viewed in any editor.)

Image Download from Cisco.com Fails with Message: "User not authorized to download file"

- Go to Configuration Manager > Tools > Security Manager Administration.
- Select Image Manager.
- Click **Test Connection** to ensure the server is reachable.
- Register your acceptance of the Cisco Encryption Software Usage Handling and Distribution Policy.

Tip The policy is found at: http://tools.cisco.com/legal/k9/controller/do/k9Check.x?eind=Y.

Image Download from Cisco.com is Slow

- Ensure proxy is configured
- Trace route from Security Manger to Cisco.com

Check for Updates Fails

Go to the Security Manager administrative settings page and test connectivity to Cisco.com. Check ***%NMSROOT%/MDC/athena/config/****-CCOMetaData.xml*** for an error in downloading metadata information for particular MDF IDs.

Message: "User not authorized to download file." Go to the Security Manager administrative settings page and test connectivity to Cisco.com. Accept the crypto agreement at http://tools.cisco.com/legal/k9/controller/do/k9Check.x?eind=Y

Image Download from External File System or Network Fileshare Fails

- Check that you have proper permissions/credentials for the external file system or fileshare.
- Open the fileshare on the client, drag the image and drop it on Image Manager.

Image Install Wizard does not Show Compatible Images

- Image Manager uses the information on Cisco.com to determine the compatibility of the images for the MDF IDs. Please download the latest images available on cisco.com by performing 'Check for Updates'. If the image is available on Cisco.com and listed as compatible for the platform, you should now be able to view these images for the device in the Image Install Wizard and also in the Compatible Images tab for the device.
- Even the Compatible Images tab for the device may not show some images that are actually compatible to the device.
- If you still don't see the images in the Install Wizard, either because you do not have Cisco.com connectivity or images are not updated on Cisco.com for that platform, you can use the drag and drop procedure to install the images on the device. You would be warned that the image is not compatible, but you can still go ahead and install the image by dragging the image and dropping it on to the device and creating a job.

Image Copy Failure –"HTTP 413 Error"

- Right-click **Device** in Image Manager > Test File Copy To Device
- Check Error Message in vmssharedsvcs.log
- If you encounter an HTTP 413 error, split the job to contain a smaller number of images in one job

Image Copy Failure –"Not enough space on disk"

- Check Device > Storage View for files on device and free space in the Image Install Location.
- If Device > Storage View shows files, then delete files from Storage to make space and retry.
- If Device > Storage View does not show files, either because it is a greenfield device or it is a Security Manager upgrade setup, rediscover only the device inventory on the device and then delete files from Storage view to make space.

Image Install Job failure –Error: "Invalid flash device"

- Check if flash exists on the device:
  - Right-click on device in IM > Test File Copy to Device
  - Connect to device, and check whether it is a multiple-context device that is being managed as a single context device in Security Manager
  - Rediscover the device selecting to discover **System Context**. Then, retry the image install job.

Image Upgrade Job on Active/Standby pair fails

- Error: ***"This host is not the 'active' device in the failover pair"***- Ensure that the failover pair is managed in Security Manager via the IP address of the active device in the failover pair and not the IP address of the standby device.
- Error: ***"Secondary device is not in standby-ready state"*** - Ensure that the devices in the failover pair are up and the standby device is in standby-ready state. Job will abort if standby device is in failed state

Image Install Job failure –Error: "SWIM1114: Device could not be reached after upgrade"

- Manually check if device is reachable. Solution: After image upgrade, device has to be re-added in Security Manager or change the Admin option to "Do not check certificate authentication"
- Check whether Tools > Admin > Device Communication > SSL Certificate Parameters > PIX/ASA/FWSM Device Authentication Certificates is set to "Retrieve while adding devices."

- After image upgrade, ensure the device has been re-added in Security Manager. Otherwise, change the Admin option to "Do not check certificate authentication."

Note You must have accepted the latest Cisco.com certificate to enable Image Manager to interface with cisco.com. You must accept the certificate from both the "Image Meta-data locator" site and the download site of the images to start downloading images successfully (see Image Manager Page, page 11-41).

No data for devices in Image Manager after Security Manager upgrade

- Rediscover the device choosing to discover only Device Inventory
- Perform a live deployment to device
- Perform an image install operation to the device

Trying to retry or rollback a job fails

- Check if any of the devices in the job are deleted from Security Manager
- Check if none of the images to be retried or rolled back to are unavailable in Security Manager. Add the images to the Security Manager repository and retry the operation

Numerics

12.1 and 12.2
- managing routers 61-3

3DES encryption algorithm
- in IKE proposals 26-6

802.1x
- 802.1x Policy page 64-5
- defining policies 64-4
- interface authorization states 64-2
- on Cisco IOS routers 64-1
- supported topologies 64-3
- understanding device roles 64-2

A

AAA
- about 48-1
- Cisco IOS routers
  - AAA Policy page 63-6
  - Accounting tab 63-10
  - Authentication tab 63-6
  - Authorization tab 63-8
  - Command Accounting dialog box 63-13
  - Command Authorization dialog box 63-10
  - defining services 63-4
  - overview 63-2
  - supported accounting types 63-3
  - supported authorization types 63-2
  - understanding method lists 63-3
- configuring access control for IPS 36-21
- configuring on firewall devices 48-1
- credentials for device access 3-4
- device administration 48-4
- local fallback 48-3
- network access 48-4
- PIX/ASA/FWSM 48-5
  - Accounting tab 48-8
  - Authentication tab 48-5
  - Authorization tab 48-7
- support 48-2
- VPN access 48-4

AAA authentication groups
- predefined 6-30

AAA firewall
- MAC exempt lists 15-26

AAA Firewall page
- Advanced Setting tab 15-20

AAA firewall policy
- advanced settings 15-20
- configuring 15-6

AAA page 15-28

AAA rules
- ACL naming conventions 12-5
- combining rules
  - example 12-27
  - interpreting results 12-25
  - procedure 12-22
- configuring AAA firewall settings (PIX/ASA/FWSM) 15-6
- configuring AuthProxy settings (IOS) 15-9
- configuring cut-through proxy (ASA) 13-23
- configuring for ASA/PIX/FWSM devices 15-4
- configuring for IOS devices 15-7
- configuring identity aware 13-21
- configuring in Map view 35-23

configuring security group aware 14-17
configuring settings
for IOS devices in Map view 35-24
for PIX/ASA/FWSM in Map view 35-24
converting IPv4 12-28
deleting 12-9
disabling 12-20
editing 12-10
enabling 12-20
managing 15-1
moving 12-19
preserving ACL names 12-4
properties 15-13
understanding 15-1
understanding how users authenticate 15-2
understanding NAT effects 12-3
understanding processing order 12-2
AAA Rules page 15-10
AAA server group objects
attributes 6-49
creating 6-48
default server groups on IOS devices 6-31
predefined authentication groups 6-30
understanding 6-27
AAA server objects
creating 6-32
HTTP-FORM settings 6-44
Kerberos settings 6-39
LDAP settings 6-40
NT settings 6-43
RADIUS settings 6-35
SDI settings 6-43
supported additional types for ASA/PIX/FWSM 6-28
supported types 6-28
TACACS+ settings 6-38
understanding 6-27
AAA servers
supported types on ASA, PIX, FWSM devices 6-28
Abort the Job dialog box 8-55
About Configuration Manager command 1-39
ABR
definition 56-75
access control list objects
creating 6-53
extended objects 6-54
standard objects 6-56
unified objects 6-58
web objects 6-57
access control lists
GET VPN security policies 29-10
policy discovery 5-14
access control lists (ACLs)
names preserved during discovery 12-4
naming conventions 12-5
resolving naming conflicts 12-7
access controls
configuring ACL names 16-23
configuring settings 16-23
configuring settings in Map view 35-24
Access Control Settings page 16-24
Access Group tab (IGMP) 55-5
Access Interface Configuration dialog box (ASA) 31-48
access permissions
Event Viewer 69-4
Health and Performance Monitor 71-3
maps 35-8
Report Manager 70-5
access policies
configuring 31-49
reference 31-45
understanding 31-44
access ports
Create and Edit Interface dialog boxes-Access Port mode 68-9
understanding 68-5
access rule
look up
from device managers 72-17

access rules
access control settings 16-24, 16-26
Access Rules page 16-10
ACL naming conventions 12-5
address requirements 16-5
Advanced dialog box 16-17
combining rules
example 12-27
interpreting results 12-25
procedure 12-22
configuring 16-7
configuring access control settings 16-23
configuring identity aware 13-21
configuring in Map view 35-23
configuring security group aware 14-17
controlling non-IP layer-2 traffic 23-1
deleting 12-9
detecting conflicts 16-28
disabling 12-20
editing 12-10
enabling 12-20
examples of event analysis
user access to server blocked 69-59
expiration dates 16-22
finding from CS-MARS events 72-45
finding from Event Viewer events 69-54
generating analysis reports 16-34
hit counts
details 16-36
how deployed 16-5
identity-aware rules
requirements 13-3
import examples 16-44
importing 16-40
IPS blocking, affect of 43-4
managing 16-1
moving 12-19
optimizing during deployment 16-46
packet tracer, analyzing with 72-23
preserving ACL names 12-4
Report Manager reports
firewall traffic reports 70-14
resolving conflicts 16-34
rule attributes 16-14
sharing ACLs among interfaces 11-18
syslog messages supported for look-up 72-46
understanding 16-1
understanding device-specific behavior 16-4
understanding global 16-3
understanding NAT effects 12-3
understanding processing order 12-2
understanding requirements when using inspection 17-4
understanding the automatic conflict detection user interface 16-30
viewing related CS-MARS events 72-42
viewing related events in Event Viewer 69-56
Accounting
Cisco IOS routers
settings 63-10
accounts and credentials
Cisco IOS routers
overview 63-14
PIX/ASA/FWSM
user accounts 51-7
user accounts, add/edit 51-7
accounts and credentials policies
Accounts and Credentials Policy page 63-16
User Accounts dialog box 63-17
ACLs
configuring names 16-23
ACS user authorization
configuring notifications when unavailable 1-27
Event Viewer 69-4
Health and Performance Monitor 71-3
how permissions affect what you can do 1-11
Report Manager 70-5
Active/Active failover

about 50-2
command replication 50-4
configuration synchronization 50-3
Active/Standby failover 50-2
Active Directory (AD)
collecting user statistics 13-25
configuring agent communication options 13-15
enabling for identity-aware firewall 13-8
identifying AD servers and agents 11-38, 13-8
requirements for identity-aware firewall 13-3
activities
accessing functions 4-8, 4-9
Activity Manager window 4-10
Approved state 4-5
approving 4-3, 4-21
benefits of 4-2
closing 4-16
creating 4-14
discarding 4-22
Edit state 4-4
locking 4-3
managing 4-1
multiple users 4-4
opening 4-15
overview 1-20
rejecting 4-21
responding to the Activity Required dialog box 4-14
states 4-4
Submitted state 4-5
submitting for approval 4-20
understanding 4-1
validating 4-18
viewing change reports 4-16
viewing status and history 4-23
working with 4-7
Activities command 1-34
Activities menu 1-36
Activity Manager window 4-10
Activity Required dialog box 4-14
Add/Edit Action Configuration dialog box 54-7
Add/Edit AnyConnect Client Image dialog box (ASA) 31-65
Add/Edit AnyConnect Custom Attributes dialog box (ASA) 31-70, 31-71
Add/Edit Applet dialog box 54-5
Add/Edit Collector dialog box 54-2
Add/Edit Content Rewrite dialog box (ASA) 31-54
Add/Edit DAP Entry Dialog Box > Device 32-30
Add/Edit File Encoding dialog box 31-55
Add/Edit Multicast Route dialog box 55-8, 55-10
description 55-9
Add/Edit PIM Neighbor Filter dialog box 55-13
Add/Edit Proxy Bypass dialog box 31-59
Add/Edit Syslog Configuration dialog box 54-7
Add AAA Rule dialog box 15-13
Add AAA Server dialog box 6-33
Add AAA Server Group dialog box 6-49
Add Access List dialog box (Allowed Hosts policy) 36-7
Add Access Rule dialog box 16-14
Add an Entry dialog box 39-30
Add AOL Class Map dialog box 17-28, 21-19
Add A Port Forwarding Entry dialog box 34-41
Add ASA Group Policies dialog box
client configuration settings 34-6
client firewall attributes 34-7
connection settings 34-33
DNS/WINS settings 34-29, 34-30
hardware client attributes 34-9
IPSec settings 34-10
overview 34-1
split tunneling settings 34-31
SSL VPN clientless settings 34-12
SSL VPN full client settings 34-19
SSL VPN settings 34-25
Technology settings 34-1
Add A Smart Tunnel Entry dialog box 34-67, 34-70
Add AS Path Entry dialog box 56-152
Add AS Path Object dialog box 56-151

Add Auto Signon Rules dialog box 34-27
Add Cat6k Block Vlan dialog box 43-16
Add Certificate dialog box 11-24
Add Certificate Filter dialog box 25-58
Add Cisco Secure Desktop Configuration dialog box 34-35
Add Client Access Rules dialog box 34-12
Add Client Update dialog box 34-81
Add Column dialog box 34-61
Add Community List Entry dialog box 56-154, 56-155
Add Community List Object dialog box 56-153
Add Custom Pane dialog box 34-62
Add Custom Signature dialog box 39-15
Add DCE/RPC Map dialog box 17-29
Add Destinations dialog box 12-11
Add Device from Network wizard
Device Credentials page 3-45
Add Devices to Group command 1-31
Add Devices to Group dialog box 3-63
Add DNS Class Map dialog box 17-28
Add DNS Map dialog box
Filtering tab 17-34, 17-35
overview 17-32
Protocol Conformance tab 17-33
Add eDonkey Class Map dialog box 17-28, 21-19
Add ESMTP Map dialog box 17-39
Add Extended Access Control Entry dialog box 6-61
Add Extended Access List dialog box 6-59
Add External Filter dialog box 21-41
Add FastTrack Class Map dialog box 17-28, 21-19
Add File Object dialog box 34-37
Add FlexConfig dialog box 7-30
Add FTP Class Map dialog box 17-28
Add FTP Map dialog box 17-42
Add Gnutella Class Map dialog box 17-28, 21-19
Add Group dialog box 3-62
Add Group Member dialog box 29-19
Add GTP Map dialog box 17-45
Add H.323 Class Map dialog box 17-28, 21-19
Add H.323 Map dialog box 17-51, 21-34
Add HSI Endpoint IP Address dialog box 17-54
Add HSI Group dialog box 17-53
Add HTTP Class Map dialog box 17-28, 21-19
Add HTTP Map dialog box 21-34
ASA 7.1.x, PIX 7.1.x, FWSM 3.x, IOS devices
Entity Length tab 17-58
Extension Request Method tab 17-61
General tab 17-57
overview 17-56
Port Misuse tab 17-62
RFC Request Method tab 17-60
Transfer Encoding tab 17-63
ASA 7.2+ and PIX 7.2+ devices 17-64
Add ICQ Class Map dialog box 17-28, 21-19
Add IKEv1 Proposal dialog box 26-10
Add IKEv2 Proposal dialog box 26-14
Add IMAP Class Map dialog box 17-28, 21-19
Add IMAP Map dialog box 21-34
Add IM Class Map dialog box 17-28
Add IM Map dialog box 21-34
ASA and PIX device 17-70
IOS device 17-73
Add Inspect/Application FW Rule wizard
Address and Port page 17-13
Inspected Protocol page 17-17
Match Traffic page 17-11
Add Inspect Parameter Map dialog box 21-31
Add Interfaces dialog box 12-13
Add IP Options Map dialog box 17-75
Add IPsec Pass Through Map dialog box 17-80
Add IPSec Transform Set dialog box 26-27
Add IPv4 Pool Object dialog box 6-92
Add IPv6 Map dialog box 17-77, 17-91
Add IPv6 Pool Object dialog box 6-93
Add Kazaa2 Class Map dialog box 17-28, 21-19
Add Key Server dialog box 29-19
Add Language dialog box 34-56
Add LDAP Attribute Map dialog box 6-46

Add LDAP Attribute Map Value dialog box 6-47
Add Link command 1-33
Add Link dialog box 35-20
Add Local Rules command 1-32
Add Local Web Filter Class Map dialog box 17-28, 21-19
Add Local Web Filter Parameter Map dialog box 21-38
Add MAC Address Pool Object dialog box 6-94
Add Map Object command 1-33
Add Map Object dialog box 35-18
Add Map Value dialog box 6-47
Add Match Condition and Action dialog box
DNS policy maps 17-36
ESMTP policy maps 17-40
FTP policy maps 17-43
GTP policy maps 17-49
H.323 (IOS) policy maps 21-35
H.323 policy maps 17-54
HTTP (Zone Based IOS) policy maps 21-35
HTTP policy maps 17-66
IM (Zone Based IOS) policy maps 21-35
IMAP policy maps 21-35
IM policy maps 17-71
IPv6 policy maps 17-78, 17-92
P2P policy maps 21-35
POP3 policy maps 21-35
SIP (IOS) policy maps 21-35
SIP policy maps 17-85, 17-95
Skinny policy maps 17-89
SMTP policy maps 21-35
Sun RPC policy maps 21-35
Web Filter policy maps 21-35
Add Match Criterion dialog box
AOL class maps 21-21
DNS class maps 17-36
eDonkey class maps 21-21
FastTrack class maps 21-21
FTP class maps 17-43
Gnutella class maps 21-21
H.323 (IOS) class maps 21-22
H.323 class maps 17-54
HTTP (IOS) class maps 21-22
HTTP class maps 17-66
ICQ class maps 21-21
IMAP class maps 21-25
IM class maps 17-71
Kazaa2 class maps 21-21
Local Web Filter class maps 21-29
MSN Messenger class maps 21-21
N2H2 class maps 21-30
POP3 class maps 21-25
SIP (IOS) class maps 21-25
SIP class maps 17-85, 17-95
SMTP class maps 21-27
Sun RPC class maps 21-29
Websense class maps 21-30
Windows Messenger class maps 21-21
Yahoo Messenger class maps 21-21
Add MSN Messenger Class Map dialog box 17-28, 21-19
Add N2H2 Parameter Map dialog box 21-39
Add N2H2 Web Filter Class Map dialog box 17-28, 21-19
Add NAT Rule dialog box
ASA 8.3+ 24-36
Add NetBIOS Map dialog box 17-81
Add Network/Host dialog box
General tab 6-83
NAT tab 24-42
Add New Device wizard
Device Credentials page 3-45
Add New Security Association dialog box 25-58
Add or Edit Plug-in Entry dialog box (ASA) 31-60
Add Other Devices dialog box 8-58
Add P2P Map dialog box 21-34
Add Permit Response dialog box 17-48
Add Per-Session NAT Rule dialog box 24-47
Add PIX/ASA/FWSM Web Filter Rule dialog box 18-5
Add PKI Enrollment dialog box
CA Information tab 26-60
Certificate Subject Name tab 26-66

Enrollment Parameters tab 26-63
overview 26-58
Trusted CA Hierarchy tab 26-67
Add Policy List Object dialog box 56-143
Add POP3 Class Map dialog box 17-28, 21-19
Add Port Forwarding List dialog box 34-40
Add Port List dialog box 6-102
Add Prefix List Entry dialog box 56-148, 56-150
Add Prefix List Object dialog box 56-146, 56-148
Add Protocol Info Parameter Map dialog box 21-33
Add Regular Expression dialog box 17-108
Add Regular Expression Group dialog box 17-108
Address Pools
PIX/ASA/FWSM 24-18
add/edit 24-19
address pools
overriding in connection profiles 30-8
Add Route Map Entry dialog box 56-137
Add Route Map Object dialog box 56-136
Add Row command 1-31
Add Rule Section dialog box 12-22
Add Server dialog box
Protocol Info Parameter maps 21-34
Add Service dialog box 6-103
Add Services dialog box 12-13
Add Single Sign On Server dialog boxes 34-42
Add SIP Class Map dialog box 17-28, 21-19
Add SIP Map dialog box 17-83, 17-93, 21-34
Add Skinny Map dialog box 17-87
Add SLA Monitor dialog box 51-10
Add Smart Tunnel Auto Signon Entry dialog box 34-72
Add Smart Tunnel Auto Signon Lists dialog box 34-71
Add Smart Tunnel Lists dialog box 34-66, 34-69
Add SMTP Class Map dialog box 17-28, 21-19
Add SMTP Map dialog box 21-34
Add SNMP Map dialog box 17-90
Add Sources dialog box 12-11
Add SSL VPN Customization dialog box 34-51
Applications 34-60
Copyright Panel 34-58
Custom Panes 34-61
Full Customization 34-59
Home Page 34-62
Informational Panel 34-57
Language 34-54
Logon Form 34-56
Logout Page 34-63
Title Panel 34-53
Toolbar 34-59
Add SSL VPN Gateway dialog box 34-64
Add Standard Access Control Entry dialog box 6-64
Add Standard Access List dialog box 6-59
Add Sun RPC Class Map dialog box 17-28, 21-19
Add Sun RPC Map dialog box 21-34
Add TCP Map dialog box 58-22
Add TCP Option Range Dialog Box 58-25
Add Text Object dialog box 7-32
Add Time Range dialog box 6-71
Add Traffic Flow dialog box 58-18
Add Transparent Firewall Rule dialog box 23-5
Add Trend Content Filter Class Map dialog box 17-28, 21-19
Add Trend Parameter Map dialog box 21-42
Add Unified Access Control Entry dialog box 6-67
Add URL Domain Name dialog box 21-45
Add URLF Glob Parameter Map dialog box 21-45
Add URL Filter Parameter Map dialog box 21-43
Add User dialog box 12-12, 36-19
Add User Group dialog box
Advanced PIX 6.3 settings 34-82
Browser Proxy settings 34-87
Client (IOS) settings 34-78
Clientless settings 34-83
Client VPN Software Update (IOS) settings 34-81
DNS/WINS settings 34-77
General settings 34-75
IOS Xauth Options settings 34-80
overview 34-73

Split Tunneling settings (Easy VPN/remote access IPSec VPN) 34-77
SSL VPN Connection settings 34-88
SSL VPN Full Tunnel settings 34-84
SSL VPN Split Tunneling settings 34-86
Technology settings 34-73
Thin Client settings 34-84
Add User Profile dialog box 43-12
Add VDI Server dialog box 34-15
Add Virtual Sensor dialog box 38-7, 38-8
Add Web Access Control Entry dialog box 6-65
Add Web Filter Map dialog box 21-47
Add WebSense Parameter Map dialog box 21-39
Add Websense Web Filter Class Map dialog box 17-28, 21-19
Add Web Type Access List dialog box 6-59
Add Windows Messenger Class Map dialog box 17-28, 21-19
Add WINS Server dialog box 34-90
Add WINS Server List dialog box 34-89
Add Yahoo Messenger Class Map dialog box 17-28, 21-19
Add Zones dialog box 12-13
admin context 59-1
administration
selecting policies to manage 5-11
administrative settings, configuring 11-1
admin password, changing 10-24
ADSL
ADSL Policy page 62-37
ADSL Settings dialog box 62-38
defining settings 62-36
supported operating modes 62-35
ADSL policies
unable to deploy 9-15
Advanced dialog box
access rules 16-17
Advanced NAT Options
PIX/ASA/FWSM
add/edit 24-29
Advanced settings
interface configuration
PIX/ASA/FWSM 46-68
AES encryption algorithm
in IKE proposals 26-6
AIM-IPS interfaces
IPS Module Interface Settings page 62-23
AIP-SSM/SSC
ASA 58-15
Alarm Indication Signal (AIS) cells 62-51
allowed hosts, configuring for IPS 36-7
Allowed Hosts policy 36-7
Analysis Engine global variables
configuring 36-30
analysis reports
generating 16-34
anomaly detection
configuring 41-6
configuring histograms 41-11
configuring learning accept mode 41-8
configuring signatures 41-4
configuring thresholds 41-11
managing 41-1
modes 41-2
understanding 41-1
understanding histograms 41-9
understanding thresholds 41-9
understanding worms 41-2
when to turn off 41-4
zones
overview 41-3
anti-spoofing 57-2
AnyConnect
client images 31-62, 31-64
profiles 31-62, 31-64
editing 31-63
AnyConnect Client Image dialog box (ASA) 31-64
AnyConnect custom attributes 31-70, 31-71
AnyConnect Profile Editor 31-63
AOL class map objects

creating 21-16
match criteria 21-21
applet
embedded event manager 54-3
Apply IPS Update command 1-35
Apply IPS Update wizard 44-7
Approve Activity command 1-36
Approve Activity dialog box 4-21
Approved activity state 4-5
Approve Deployment Job dialog box 8-20, 8-39
Area Border Router
See ABR 56-75
ARP
PIX/ASA/FWSM
configuration 47-5
inspection 47-5
inspection, enable/disable 47-6
table 47-3
ARP table
static entry 47-3, 47-5
ASA
ASDM 72-15
CX 58-17
Auth Proxy Configuration 58-17
CX module
detecting 72-21
Failover
Add Failover Group 50-25
edit bridge group 50-17
FirePOWER module
detecting 72-21
IPS, QoS, and Connection Rules
ASA CX Auth Proxy Configuration 58-17
IPS modules 58-15
policy discovery 5-13
rollback, commands to recover from failover misconfiguration 8-68
rollback command conflicts 8-67
rollback restrictions for failover devices 8-65
rollback restrictions for multiple context mode 8-64
security contexts
allocate interfaces 59-12
configuration 59-9
viewing allocated interfaces 59-12
setting up AUS or CNS 2-8
setting up SSL (HTTPS) 2-3
TCP State Bypass 58-3
ASA 5505
Management IPv6 47-11
ports and interfaces 46-6
ASA 8.3+
NAT policies
Add/Edit NAT rules dialog boxes 24-36
Translation Rules page 24-34
ASA Cluster Load Balance page 31-5
ASA CX
CX
about 58-17
ASA devices
5505
hardware port configuration 46-61
AAA support 6-28
about 46-1
adding or changing modules 3-40
adding SSL thumbprints manually 9-5
Bridge Groups
add/edit 46-62
Catalyst Service Module 46-1
changing those selected for reports 70-22
configuring for event management 69-28
configuring for report management 70-3
configuring IKE and IPsec policies 26-1
configuring IKEv2 authentication 26-68
configuring transparent firewall rules 23-1
Easy VPNs
connection profiles 28-13
Event Viewer support 69-4
FlexConfig object samples 7-20

global access rules 16-3
identity-aware services
configuring to provide 13-7, 14-8
interfaces 46-26
add/edit 46-31
Advanced tab 46-41
configuring 46-3
edit EtherChannel-assigned interface 46-12
EtherChannels 46-9, 46-13
General tab 46-33
IP Type 46-58
IPv6 46-47, 46-73
IPv6, add/edit 46-52
IPv6, add/edit prefixes 46-54, 46-56
LACP 46-12
MAC address 46-60
PPPoE Users 46-71
VPDN groups 46-72
licenses 2-9
monitoring service level agreements 51-8
object group search 16-25
packet capture, using 72-30
packet tracer, using 72-23
remote access SSL VPNs
advanced settings 31-72
Anyconnect client settings 31-62, 31-64
browser plug-ins 31-60
configuring HTTP/HTTPS proxies and proxy bypass 31-57
content rewrite rules 31-53
encoding rules 31-55
Kerberos Constrained Delegation (KCD) 31-66, 31-69
other settings 31-51
performance settings 31-52
server certificate verification settings 31-30, 31-32, 31-73
shared license 31-74
shared license clients (ASA) 31-76
shared license servers (ASA) 31-77
remote access VPNs
access policies (ASA), configuring 31-49
access policies (ASA), reference 31-45
access policies (ASA), understanding 31-44
AnyConnect client image settings (ASA) 31-65
AnyConnect custom attributes (ASA) 31-70, 31-71
certificate to connection profile map policy (IKEv1) 31-36
certificate to connection profile map rules (IKEv1 IPSec) 31-37
cluster load balancing 31-5
configuring bookmarks 31-82
configuring portal appearance 31-78
configuring WINS servers for file system access 31-88
connection profiles 31-7, 31-8
creating IPSec 30-25
creating SSL 30-14
customizing 31-77
device support 30-8
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
fragmentation settings 26-31, 26-44
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
IKE proposals 26-9
IKEv2 settings 26-37
IPsec proposals 31-41
ISAKMP/IPsec settings 26-33
managing 31-1
NAT settings 26-42
policy overview 31-2
post URL method and macro substitutions in bookmarks 31-84
proxy bypass rules (ASA) 31-59
Public Key Infrastructure (PKI) 26-56
secure desktop manager policies 32-9
smart tunnels 31-85

understanding IKE 26-5
understanding NAT settings 26-41
wizard 30-13
Report Manager reports
firewall summary botnet reports 70-15
firewall traffic reports 70-14
general VPN reports 70-16
VPN top reports 70-16
selecting for Event Viewer 69-34
selecting policy types to manage 5-11
SSL certificate configuration 11-22
ASA group policies objects
client configuration settings 34-6
client firewall attributes 34-7
connection settings 34-33
DNS/WINS settings 34-29, 34-30
hardware client attributes 34-9
IPSec settings 34-10
split tunneling settings 34-31
SSL VPN clientless settings 34-12
SSL VPN full client settings 34-19
SSL VPN settings 34-25
technology settings 34-1
ASA Image Management 73-16, 73-32
ASAv
about 46-1
ASBR
definition 56-75
ASCII limitations for text 1-50
ASDM
access rule look-up 72-18
device manager 72-15
AS path objects
properties 56-151
ASR
zone-based firewall
global parameters 21-50
restrictions 21-3
assignment overview 1-20
Assignments tab, Policy view 5-54
Assign Shared Policy command 1-32
Assign Shared Policy dialog box 5-44
Asymmetric Digital Subscriber Line (ADSL)
on Cisco IOS routers 62-34
Asymmetric Routing Groups 46-6
Asynchronous Transfer Mode (ATM) 62-47
ATM 62-47
virtual channel connections (VCCs) 62-47
virtual channel identifier (VCI) 62-47
virtual path connections (VPCs) 62-47
virtual path identifier (VPI) 62-47
Attack Response Controller 43-1
attacks
broadcast 17-4
Denial of Service (DoS) 17-5
spoofing 17-4
SYN flooding 17-5
audit logs
configuring default settings 11-62
purging entries 10-23
understanding 10-19
working with 10-19
Audit Message Detail dialog box 10-21
Audit Report command 1-34
audit reports
generating and viewing 10-20
understanding 10-19
working with 10-19
Audit Report window 10-21
AUS
deploying configurations 8-41
deployment method 8-10
setting up 2-8
setting up on PIX Firewall and ASA devices 2-8
Authentication
Cisco IOS routers
settings 63-6
authentication

routing protocols 56-75
Authentication-Authorization-Accounting
see AAA 48-1
Authentication Header (AH) encryption algorithm 26-30
authentication methods
certificates (RSA signatures) 26-8
in IKE proposals 26-8
preshared keys 26-8
authentication testing
SSH 2-6
Authorization
Cisco IOS routers
settings 63-8
authorization proxy (AuthProxy)
configuring AAA rules 15-7
AuthProxy
configuring settings in Map view 35-24
Auth Proxy Configuration
ASA CX 58-17
AuthProxy dialog box 15-19
AuthProxy settings policy
configuring 15-9
autolink
omitting reserved networks from maps 11-3
automatic conflict detection
resolving conflicts 16-34
understanding 16-28
understanding the user interface 16-30
using 16-28
autonomous system paths
See AS paths
auto signon rules
ASA group policy objects 34-27
Auto Update Server (AUS)
adding 3-36
licensing 10-17
PIX/ASA/FWSM 52-1
add/edit server 52-3, 53-2, 53-3, 53-4
troubleshooting deployment 9-18
Auto Update Server Properties dialog box 3-38
Available Bit Rate (ABR) 62-48
Available Servers dialog box 3-39

B

background image, map
deleting 35-13
importing 35-13
scale and position 35-13
setting 35-13
backup
event data store 69-36
backup.pl command 10-25
Backup command 1-35
backups, Security Manager database 10-25
bandwidth
VPN user reports 70-16
banners
configuring on firewall devices 48-9
benefits of product 1-2
BGP routing
BGP Routing Policy page 67-4
defining routes 67-2
Neighbors dialog box 67-6
on Cisco IOS routers 67-1
PIX/ASA/FWSM 56-2, 56-3
General tab 56-5
IPv4 Family - Aggregate Address configuration 56-9, 56-22
IPv4 Family - Filter configuration 56-10
IPv4 Family - General tab 56-7, 56-21
IPv4 Family - Neighbor configuration 56-11, 56-24
IPv4 Family - Network configuration 56-17, 56-29
IPv4 Family - Redistribution configuration 56-18, 56-30
IPv4 Family - Route Injection configuration 56-19, 56-31
IPv4 Family tab 56-6, 56-20

redistributing routes 67-3
Redistribution Mapping dialog box 67-7
Redistribution tab 67-7
Setup tab 67-5
Bidirectional Neighbor Filter 55-14
Bidirectional Neighbor Filter tab
PIM 55-13
blocking, IPS
configuring 43-7
configuring ARC 43-1
configuring blocking devices 43-14
configuring master blocking sensors 43-13
configuring never block hosts and networks 43-17
configuring router blocking interfaces 43-15
configuring user profiles 43-12
configuring VLAN blocking interfaces 43-16
general options 43-10
master blocking sensor 43-6
policy 43-8
rate limiting 43-4
router and switch blocking devices 43-4
strategies 43-3
understanding 43-1
Blocking page 43-8
Boot image/configuration
PIX/ASA 48-10
add/edit 48-12
bootstrap configuration
Failover 50-26
Botnet Traffic Filter Drop Rules Editor 19-13
botnet traffic filter rules
adding static entries 19-5
blocking blacklisted traffic 19-6
configuring DNS snooping 17-19
configuring in Map view 35-23
configuring the dynamic database 19-4
configuring with IPS global correlation 42-1
databases 19-1
Device Blacklist dialog box 19-15
Device Whitelist dialog box 19-15
Drop Rules Editor 19-13
Dynamic Blacklist Configuration tab 19-10
enabling DNS snooping 19-6
field definitions 19-9
illustrations 19-1
mitigating botnet activity 69-65
monitoring
activity using ASDM 69-64
activity using Event Viewer 69-62, 69-64
overview 69-61
understanding botnet syslog events 69-61
overview 19-1
preserving ACL names 12-4
Report Manager reports
firewall summary botnet reports 70-15
task flow 19-2
traffic classification 19-6
Traffic Classification dialog box 19-12
Traffic Classification tab 19-11
understanding 19-1
understanding NAT effects 12-3
understanding processing order 12-2
Whitelist/Blacklist tab 19-14
bridge group
failover
editing 50-17
Bridge Groups
ASA/FWSM
add/edit 46-62
bridge groups
defining 63-20
FWSM 3.1 47-3
Bridging
ASA 5505
Management IPv6 47-11
PIX/ASA/FWSM
ARP configuration 47-5
ARP Inspection 47-5

ARP Inspection, enable/disable 47-6
ARP Table 47-3
MAC Address, add/edit 47-8
MAC Address Table 47-8
MAC Learning 47-9
MAC Learning, enable/disable 47-9
Management IP address 47-10
bridging
Cisco IOS routers
Bridge Group dialog box 63-21
Bridging Policy page 63-21
BVI interfaces 63-19
overview 63-18
configuring transparent firewall rules 23-1
PIX/ASA/FWSM
about 47-1
configuring on 47-1
broadcast attacks, preventing 17-4
broadcasts
enabling directed on routers 62-20
browser plug-ins
configuring 31-60
Bundles 73-13
bypass mode
configuring for IPS 37-12

C

CA server authentication methods
SCEP (Simple Certificate Enrollment Protocol) 26-51
Cat6k Device dialog box 43-14
Catalyst 6500/7600 devices
configuring FWSM in site-to-site VPNs 25-47
configuring SSH 2-6
default transport protocol 11-22
deployment 8-28
FlexConfig object samples 7-22
IPS blocking devices 43-4
policy discovery for FWSM 5-13
rollback restrictions 8-65
Service Modules 46-1
Catalyst 6500/7600 switches
including in deployment jobs 8-28
Catalyst devices
policy discovery 5-13
remote access VPNs
Dynamic VTI/VRF Aware IPsec settings 33-7
high availability 33-11
IPsec proposals 33-4
user group policies 33-13
VPNSM/VPN SPA/VSPA settings 33-6
Catalyst platform policies
IDSM settings policy
Create and Edit IDSM Data Port VLANs dialog boxes 68-49
Create and Edit IDSM EtherChannel VLANs dialog boxes 68-49
IDSM Settings page 68-47
IDSM Slot-Port Selector dialog box 68-50
interfaces/VLANs policy
Access Port Selector dialog box 68-30
Create and Edit Interface dialog boxes-Access Port mode 68-9
Create and Edit Interface dialog boxes-Dynamic Port mode 68-18
Create and Edit Interface dialog boxes-Other mode 68-24
Create and Edit Interface dialog boxes-Routed Port mode 68-12
Create and Edit Interface dialog boxes-subinterfaces 68-22
Create and Edit Interface dialog boxes-Trunk Port mode 68-14
Create and Edit VLAN dialog boxes 68-28
Create and Edit VLAN Group dialog boxes 68-34
Interfaces tab 68-8
Service Module Slot Selector dialog box 68-35
Summary tab 68-3
Trunk Port Selector dialog box 68-31
VLAN Groups tab 68-33

VLAN Selector dialog box 68-35
VLANs tab 68-27
VLAN access lists policy
Create and Edit VLAN ACL Content dialog boxes 68-41
Create and Edit VLAN ACL dialog boxes 68-41
VLAN Access Lists page 68-39
Catalyst Summary Info command 1-35
Catalyst switches
configuring SSH 2-6
default transport protocol 11-22
showing modules, security contexts, and virtual sensors 3-56
Catalyst switches/7600 routers
troubleshooting deployment 9-16
Catalyst switches and 7600 devices
IDSM mode support 68-43
interface deployment failure 9-16
internal VLAN deployment failure 9-16
supported VTP modes 68-1
Catalyst switches and 7600 Series routers
access ports 68-5
Catalyst Summary Info page 68-2
defining IDSM Data Port VLANs 68-46
defining IDSM EtherChannel VLANs 68-44
defining ports 68-6
defining VACLs 68-37
defining VLAN groups 68-32
defining VLANs 68-26
deleting IDSM Data Port VLANs 68-47
deleting IDSM EtherChannel VLANs 68-45
deleting ports 68-7
deleting VACLs 68-38
deleting VLAN groups 68-33
deleting VLANs 68-27
discovering policies 68-1
generating interface names 68-6
IDSM settings 68-43
IDSM Settings page 68-47
interfaces 68-5
managing 68-1
routed ports 68-5
trunk ports 68-5
viewing interface and VLAN summary 68-3
VLAN Access Lists page 68-39
VLAN ACLs (VACLs) 68-36
VLAN groups 68-31
VLANs 68-25
Catalyst VPN Service Port Adapters (VSPAs)
configuring 25-42
Catalyst VPN Services Module (VPNSM)
configuring 25-42
configuring in remote access VPNs 33-6
Catalyst VPN Shared Port Adapter (VPN SPA)
configuring 25-42
configuring in remote access VPNs 33-6
categories
using 6-13
cautions
significance of 2-lxiii
CCO settings 11-4
CDP
configuring mode for IPS 37-12
CEF Interface Settings dialog box 62-27
CEF interface settings policies 62-25
certificates
accepting 11-4, 11-52
retrieving 11-4, 11-52
viewing 11-4, 11-52
certificates, SSL
adding thumbprints manually 9-5
configuring default settings for how handled 11-22
managing IPS 44-10
certificates for ASA image downloads 11-4
certificates for IPS package downloads 11-52
certificate to connection profile map policies
configuring policy 31-36
configuring rules 31-37

certificate trust management 11-4, 11-52
Change Report dialog box 4-18
change reports
  selecting session in non-Workflow mode 4-18
  viewing 4-16
Change Reports command 1-34
Checkpoint migration
  configuring object group search on ASA 8.3+ devices 16-25
Choose a file dialog box 34-39
Cisco 7600 Series routers
  managing 68-1
Cisco AnyConnect Profile Editor 31-63
Cisco Configuration Engine
  troubleshooting device setup and deployment 9-18
Cisco Discovery Protocol (CDP)
  enabling CDP on router interfaces 62-18
Cisco Express Forwarding (CEF)
  CEF Interface Settings policy 62-26
  CEF router interface settings policies 62-25
  importance for QoS 66-2
Cisco IOS IPS
  affect of load balancing 45-8
  configuration files 45-3
  configuration overview 45-4
  configuring 45-1
  configuring general settings 45-7
  configuring interface rules 45-9
  getting started 36-1
  initial preparation of router 45-5
  lightweight signature engines 45-2
  limitations and restrictions 45-3
  selecting signature category 45-6
  understanding 45-1
  understanding subsystems and revisions 45-2
Cisco IOS Routers
  configuring IOS IPS 45-1
  IPS blocking devices 43-4
Cisco IOS routers
802.1x 64-1
AAA 63-2
accounts and credentials 63-14
ADSL 62-34
advanced interface settings 62-13
available interface types 62-2
basic interface settings 62-1
BGP routing 67-1
configuring SSH 2-6
CPU settings 63-25
default AAA server groups 6-31
deploying configurations using TMS 8-43
dialer interfaces 62-28
discovering policies 61-3
Domain Name System (DNS) 63-74
Dynamic Host Configuration Protocol (DHCP) 63-87
EIGRP routing 67-8
host and domain names 63-77
HTTP 63-28
interface deployment failure 9-14
IOS 12.1 and 12.2 61-3
licenses 2-10
line access 63-35
managing 61-1
memory settings 63-78
NAT 24-5
  designating interfaces 24-6
  dynamic rules 24-10
  static rules 24-6
  timeouts 24-13
NetFlow 65-1, 65-5, 65-12
Network Admission Control (NAC) 64-8
Network Time Protocol (NTP) 63-96
optional SSH settings 63-63
OSPF routing 67-19
permanent virtual connections (PVCs) 62-47
platform policies 61-1
Point-to-Point Protocol (PPP) 62-71
policy discovery 5-13

quality of service (QoS) 66-1
RIP routing 67-42
Secure Device Provisioning (SDP) 63-81
setting up SSL (HTTPS) 2-4
SHDSL 62-41
SNMP 63-66
static routing 67-50
syslog logging 65-1
time zone settings 63-22
transparent bridging 63-18
Cisco IOS Software
FlexConfig object samples 7-22
selecting policy types to manage 5-11
Cisco Prime Security Manager
see PRSM 72-20, 72-21
Cisco Secure Desktop configuration objects
creating 33-18
Cisco Security Management Suite server
logging into or exiting 1-12
Cisco Technical Assistance Center
creating diagnostic file 10-28
generating data 10-28
generating deployment or discovery status reports 10-30
generating partial database backup 10-30
Cisco Trust Agent (CTA) 64-9
CiscoWorks Common Services
backing up and restoring Security Manager 10-24
logging into or exiting 1-12
CiscoWorks user authorization, affect on what you can do 1-11
Class-Based Policing 66-6
class maps
understanding 6-78
Clear Connection Configuration dialog box 15-25
clear xlate
PIX/ASA/FWSM platform 60-1
CLI commands
FlexConfig objects 7-2
client applications 72-2
client connection characteristics
configuration modes 28-3
configuring policies for Easy VPN 28-7
extended authentication (xauth) 28-4
clientless access mode 30-4
client settings
configuring AnyConnect 31-64
understanding AnyConnect 31-62
client-side file browsing 1-50
enabling or disabling 11-10
CLI prompt
configuring on firewall devices 48-12
Clock
PIX/ASA/FWSM 48-14
clock
Cisco IOS routers
overview 63-22
clock settings
Cisco IOS routers
Clock Policy page 63-23
Clone Device command 1-30
Clone Policy Bundle dialog box 5-58
Clone Policy command 1-32
Clone Policy dialog box 5-47
Close Activity command 1-36
Close All Reports command (Report Manager) 70-8
Close Report command (Report Manager) 70-8
Close Ticket command 1-37
cluster, server
managing 10-2
overview 10-2
splitting server 10-3
synchronizing shared policies 10-5
Cluster Information page, device properties 3-50
clustering 3-9
cluster load balancing
configuring 31-5
understanding 31-5

understanding FQDN redirection 31-5
CNS
deploying configurations 8-41
deployment method 8-10
setting up on PIX Firewall and ASA devices 2-8
color rules, configuring in Event Viewer 69-40
Combine Rules Selection Summary dialog box 12-24
commands
Activities menu 1-36
Edit menu (Configuration Manager) 1-31
Event Viewer File menu 69-9
Event Viewer View menu 69-10
File menu (Configuration Manager) 1-30
Help menu (Configuration Manager) 1-38
Launch menu 1-37
Manage menu 1-34
Map menu 1-33
Policy menu (Configuration Manager) 1-32
Report Manager menus 70-8
Tickets menu 1-36
Tools menu (Configuration Manager) 1-34
View menu (Configuration Manager) 1-31
Common Services
licensing 10-17
communication, device
troubleshooting 9-8
community list objects
properties 56-153
configurable dashboard for IPS and FW 72-1
configuration
initial Security Manager 1-25
understanding rollback 8-63
Configuration Archive
adding configurations from devices 8-59
overview 8-15
rolling back to archived configuration files 8-70
rolling back when deploying to file 8-71
settings 11-6
version viewer 8-60
viewing and comparing configuration versions 8-59
viewing transcripts 8-62
window 8-23
Configuration Archive command 1-34
Configuration Archive page 11-6
Configuration Engine
adding 3-36
setting up 2-8
Configuration Engine Properties dialog box 3-38
configuration files
deploying in non-Workflow mode 8-28
deploying in Workflow mode 8-34, 8-39
deploying to 8-11
deploying to an AUS or CNS 8-41
deploying to a TMS 8-43
deployment process overview 8-1
factory-default configurations 46-2
previewing 8-44
redeploying to devices 8-53
rolling back after deploying to file 8-71
rolling back to archived configurations 8-70
rolling back to devices 8-69
selecting 1-50
web VPN policy discovery restrictions 3-8
configuration location, configuring for IOS IPS 45-7
Configuration Manager
overview 1-14
using 1-14
configurations
adding to the Configuration Archive 8-59
avoiding out-of-band changes 8-47
detecting out-of-band changes 8-45
rollback, commands to recover from failover misconfiguration 8-68
rollback command conflicts 8-67
rolling back 8-63
rolling back Catalyst 6500/7600 8-65
rolling back failover devices 8-65
rolling back IPS and IOS IPS 8-66

rolling back multiple context mode 8-64
understanding out-of-band changes 8-12
viewing and comparing 8-59
configuration session
selecting session for change reports 4-18
viewing change reports 4-16
configuration sessions
discarding 4-22
configuration views 1-14
Configure dialog box 17-22
Configure DNS dialog box 17-19
Configure ESMTP dialog box 17-20
Configure Fragments dialog box 17-20
Configure Hardware Ports
ASA 5505 46-61
Configure IMAP dialog box 17-21
Configure POP3 dialog box 17-21
Configure RPC dialog box 17-21
Configure SMTP dialog box 17-20
Config Version Viewer (Preview Configuration) dialog box 8-44
conflict analysis reports
generating 16-34
conflict detection
resolving conflicts 16-34
understanding 16-28
understanding the user interface 16-30
using 16-28
connection
PIX/ASA/FWSM
identity-aware rules 13-21
rules 58-5
Connection Alias dialog box 31-25, 31-34
Connection Profile dialog box
AAA tab 31-13
General tab 31-10
IPSec tab 31-19
Secondary AAA tab 31-17
SSL tab 31-22
connection profiles
configuring 31-7
configuring for Easy VPN 28-13
properties
AAA 31-13
general 31-10
IPSec 31-19
policy overview 31-8
secondary AAA 31-17
SSL 31-22
sharing among multiple ASAs 30-8
Connection Profiles page 31-8
Connection Settings
MPC rule wizard
tab 58-8
connection timeout
device communication settings 11-22
Connection URL dialog box 31-25
connectivity, testing device 9-1
console
Cisco IOS routers
AAA tab 63-44
Accounting tab 63-47
Authentication tab 63-44
Authorization tab 63-45
Console Policy page 63-42
Setup tab 63-42
console port
Cisco IOS routers
defining AAA settings 63-37
defining setup parameters 63-35
Console timeout
PIX/ASA/FWSM 49-1
Constant Bit Rate (CBR) 62-48
contained modules
showing 3-56
content rewrite rules
defining for SSL VPN on ASA 31-53
Context-Based Access Control

choosing interfaces 17-2
configuring 17-5
configuring identity aware 13-21
preventing DoS attacks on IOS devices 17-5
selecting protocols 17-3
understanding 17-2
understanding access rule requirements 17-4
Context Editor dialog box (IOS) 33-15
contexts
see “security contexts” 59-1
continuity check (CC) cells 62-51
control plane (CP)
defining QoS on 66-12
policing on 66-9
Control Plane Policing 66-9
conventions 2-lxiii
cookie challenges 26-37
Copy command 1-31, 12-9
Copy Policies Between Devices command 1-32
Copy Policies wizard 5-33
CPU settings
defining utilization settings 63-25
overview 63-25
CPU Throttling Policy 11-33
CPU utilization
CPU Policy page 63-26
Create a Clone of Device dialog box 3-56
Create Activity dialog box 4-14
Create a Policy dialog box 5-54
Create Discovery Task dialog box 5-18
Create Extranet VPN Topology wizard
overview 25-66
Create Filter dialog box 1-46
Create Group Policy wizard
Clientless and Thin Client Access Modes page 30-24
Full Tunnel page 30-21
Group Policy page 30-20
using 30-20
Create Overrides for Device dialog box 6-20
Create Policy Bundle dialog box 5-57
Create Text Object dialog box 7-32
Create Ticket dialog box 4-14
Create VPN Topology wizard
Device Selection page 25-32
Edit Endpoints dialog box 25-34
Endpoints page 25-34
GET VPN Group Encryption page 25-54
GET VPN Peers page 25-60
High Availability page 25-52
Name and Technology page 25-30
overview 25-28
VPN Defaults page 25-62
credential objects
attributes 28-9
credentials
configuring on firewall devices 48-17
device manager validation 72-14
IPS module 3-20
service module 3-19
testing 9-1
understanding device 3-4
Credentials page
HTTPS port number
overriding with HTTP policy 3-47
Credentials page, device properties 3-45
crypto maps
understanding 26-19
CSC
MPC rule wizard
tab 58-8
CSDM Policy Editor dialog box 32-46
CS-MARS
access to Security Manager 72-37
configuring servers 11-7
discovering or changing controller used by device 72-39
events
historical and real-time lookup 72-41

looking up 72-41
integrating with Security Manager 72-36
integration with Security Manager 72-36
looking up Security Manager policies based on events 72-45
NetFlow 72-47
query
troubleshooting 72-40
registering in Security Manager 72-38
supported log messages 72-46
viewing access rule events 72-42
viewing IPS signature events 72-44
CS-MARS page 11-7
CSMDiagnostics.zip
setting debug options 11-11
CSMDiagnostics.zip file, creating 10-28
CSM Mobile 72-11
settings page 11-9
CSM Monitor widget 72-7
CSM tab, Licensing page 11-57
CSV (comma-separated values) files
supported formats for device inventory 10-9
CSV file
export HPM data as 71-31
Customize Desktop Settings page 11-10
Customized Toolbar command 1-32
Custom Protocol dialog box 17-22
Custom Report List command (Report Manager) 70-9
Cut command 1-31, 12-9
cut-through proxy, configuring 13-23
CX
ASA module
detecting 72-21
CXSC
MPC rule wizard
tab 58-8

D

Dashboard
CSM Mobile settings page 11-9
Dashboard tabs
default view 72-8
re-arranging 72-8
Dashboard widgets for device heath trends 72-2
database
backing up 10-25
backing up and restoring 10-24
generating partial backups for TAC 10-30
restoring 10-27
DCE/RPC policy map objects
creating 17-22
properties 17-29
DCS.properties file
DCS.doSerialAccessForFWSMVCs property 9-17
DCS.FWSM.checkThreshold property 9-17
SSH settings 9-7
warning message expression properties 9-10
DDNS
PIX/ASA/FWSM 52-18
add interface rules 52-19
update methods 52-19
update methods, add/edit 52-20
dead-peer detection (DPD) 26-33
debugging
configuring debug levels 11-11
Debug Options page 11-11
Default Report Settings command (Report Manager) 70-9
defaults, configuring 11-1
Delete Device command 1-30
Delete Map command 1-33
Delete Map dialog box 35-10
Delete Row command 1-31
Denial of Service (DoS)
preventing in SMTP using zone based firewall 21-27
denial of service (DoS)

preventing using unicast reverse path forwarding (RFP) 62-20
Denial of Service (DoS) attacks
configuring inspection settings to mitigate 17-111
preventing on IOS devices using inspection 17-5
denial of service (DoS) attacks
preventing using IKEv2 cookie challenge 26-37
deny
inspection
rules 17-5
Deploy command 1-31
Deploy Job dialog box 8-39
deployment
Add Other Devices dialog box 8-58
Auto Update Server 8-41
Catalyst 6500/7600 devices 8-28
changes not deployed when using schedules 8-55
changing device message severity level to ignore errors 9-10
changing FWSM multiple-context deployment to serial 9-17
Cisco Networking Services configuration engine 8-41
clearing XLATE on 60-1
configuration files, to 8-11
configurations 8-28
creating jobs in Workflow mode 8-35
creating or editing schedules 8-55
Deployment Manager window 8-16
device communication settings 9-4
devices, directly to 8-9
devices, through intermediate server 8-10
Edit Deploy Method dialog box 8-30
Edit Selected Deployment Method dialog box 8-30
error attempting to remove unreferenced object 9-12
errors
OS version mismatches 8-13
generating status report 10-30
handling OS version mismatches 8-13
managing 8-1
methods 8-8
minimum memory errors for ASA 8.3+ 9-12
non-Workflow mode 8-3
optimizing access rules 16-46
out-of-band changes
avoiding 8-47
detecting and analyzing 8-45
understanding 8-12
process overview 8-1
rolling back archived configurations 8-70
rolling back configurations 8-63
rolling back configurations, Catalyst 6500/7600 8-65
rolling back configurations, command conflicts 8-67
rolling back configurations, commands to recover from failover misconfiguration 8-68
rolling back configurations, failover devices 8-65
rolling back configurations, IPS and IOS IPS devices 8-66
rolling back configurations, multiple context mode 8-64
rolling back configuration when deploying to file 8-71
rolling back to last deployed configuration 8-69
setting debug options 11-11
suspending or resuming schedules 8-58
system settings 11-13
task flow
non-Workflow mode 8-3
Workflow mode 8-5
tips for successful jobs 8-27
TMS server 8-43
troubleshooting 9-1, 9-9
ADSL or PVC deployment failures 9-15
AUS problems 9-18
Catalyst interface settings 9-16
Catalyst internal VLANs 9-16
Catalyst switch and modules 9-16
Configuration Engine problems 9-18
Error Writing to Server messages 9-15
HTTP Response Code 500 messages 9-15

layer 2 interfaces 9-15
mixing deployment methods with routers and VPNs 9-14
router interface settings 9-14
routers 9-14
Security Manager cannot contact device 9-12
VPNs with routing processes 9-13
troubleshooting device communication 9-8
troubleshooting router connection failures 2-2
troubleshooting SSL certificate errors 9-5
troubleshooting VRF-aware IPsec on Catalyst 6500/7600 devices 25-17
understanding 8-1
understanding configuration rollback 8-63
using a Cisco Networking Services (CNS) server 8-41
viewing device details 8-26
viewing job summary 8-26
viewing status and history for jobs and schedules 8-26
viewing transcripts 8-62
Warning - Partial VPN Deployment dialog box 8-31
Workflow mode 8-5, 8-34, 8-39
working with 8-25
Deployment—Create or Edit a Job dialog box 8-35
deployment jobs
aborting 8-55
approval 8-7
approving 8-39
creating and editing in non-Workflow mode 8-28
creating and editing in Workflow mode 8-35
Deployment Manager 8-16
discarding 8-41
including devices in 8-8
multiple users 8-8
redeploying 8-53
rejecting 8-39
states
non-Workflow mode 8-4
Workflow mode 8-6
submitting 8-38
viewing history 8-26
Deployment Manager
overview 8-15, 8-16
Deployment Manager window 8-16
Deployment Schedules tab 8-21
Deployment page
PIX/ASA/FWSM Platform
clear xlate 60-1
Deployment Schedules tab 8-21
Deployments command 1-34
Deployment Settings page 11-13
Deployment Status Details dialog box 8-32
Deployment Workflow Commentary dialog boxes 8-20
Deploy Saved Changes dialog box 8-28
DES encryption algorithm
in IKE proposals 26-6
Designated Router
PIX/ASA/FWSM 55-12
Destination Contents dialog box 12-14
Dest Port Map dialog box 41-12
Detect Out of Band Changes command 1-35
device
AAA administration 48-4
firewall types 46-1
viewing inventory status 72-12
Device Access
FWSM
Resources, add/edit 51-4
PIX/ASA/FWSM 49-1
console timeout 49-1
host name 51-1
HTTP configuration 49-3
HTTP page 49-2
ICMP rules 49-4
ICMP rules, add/edit 49-5
Management Access interface 49-6
Secure Shell, add/edit host 49-8
Secure Shell (SSH) 49-7, 49-8

Server Access 52-1, 53-1
SNMP host access 49-22
SNMP page 49-17
SNMP Trap configuration 49-19
Telnet configuration 49-29
Telnet page 49-29
user accounts 51-7
user accounts, add/edit 51-7
device access policies
defining 63-14
Device Admin
FWSM
Resources 51-3
device administration policies
configuring on firewall devices 48-1
device authentication
adding SSL thumbprints manually 9-5
SSL certificate default configuration 11-22
Device Blacklist dialog box 19-15
device clusters 3-9
device communication
changing device message severity level 9-10
managing settings 9-4
routers without K8/K9 crypto image 9-8
Security Manager cannot contact device after deployment 9-12
troubleshooting failures 9-8
Device Communication page 11-21
device communications
troubleshooting 9-1
device communication settings
connection timeout 11-22
retry count 11-22
socket read timeout 11-22
Device Connectivity Test dialog box 9-3
device credentials
understanding 3-4
Device Credentials page 3-45
Device Delete Validation dialog box 3-59
device groups 3-59, 3-62
adding or removing devices 3-63
creating group types 3-62
deleting groups or types 3-63
understanding 3-60
Device Groups page 3-49, 11-24
device health trends in Dashboard 72-2
Device Information page - Add Device from File 3-33
Device Information page - Configuration File 3-23
Device Information page - Network 3-14
Device Information page- New Device 3-27
device inventory
exporting
DCR, CS-MARS, Security Manager formats 10-6
device with policies 10-6
overview 10-6
supported CSV formats 10-9
using command line utility 10-10
importing
device with policies 10-13
importing with policies 10-13
managing 3-1
sharing with PRSM 72-22
testing device connectivity 9-1
understanding 3-1
understanding contents 3-3
understanding device clusters 3-9
understanding generic devices 3-8
working with 3-36
device manager
access rule look up 72-17
ASDM 72-15
access rule look-up 72-18
credentials 72-14
IDM 72-15
PDM 72-15
prerequisites 72-16
SDM 72-16

access rule look-up 72-19
starting from HPM 71-3, 71-27
starting from Security Manager 72-14
troubleshooting 72-16
xdm-launcher.exe 72-16
Device Manager command 1-37
Device Properties
Cluster Information page 3-50
Credentials page 3-45
Device Groups page 3-49
General page 3-41
Policy Object Override pages
general reference 3-52
device properties
changes with policy effects 3-54
changing critical 3-52
image version changes with no policy effects 3-53
understanding 3-6
viewing or changing 3-40
Device Properties command 1-35
Device Properties page
creating object overrides 6-19
deleting overrides 6-21
overview 3-40
device response
to appear as an error message 9-10
devices
adding 3-6
adding configurations to the Configuration Archive 8-59
adding from configuration files 3-22
adding from inventory file 3-31
adding from network 3-12
adding local rules to shared policies 5-45
adding manually 3-26
adding or changing modules 3-40
assigning shared policies 5-44
avoiding out-of-band changes 8-47
changing critical properties 3-52
changing those selected for reports 70-22
cloning or duplicating 3-56
cloning shared policies 5-47
communication requirements 2-1
communication settings and certificates 9-4
configuring ASA licenses 2-9
configuring IOS licenses 2-10
configuring local policies 5-31
copying policies between 5-33
creating policy object overrides 6-19
deleting from inventory 3-58
deleting policy object overrides 6-21
deployment through intermediate server 8-10
deployment to 8-9
detecting out-of-band changes 8-45
discovering or changing CS-MARS controller 72-39
discovering policies 5-12
discovering policies on existing devices 5-15
dynamic IP addresses 3-36
image version changes with no policy effects 3-53
including in deployment jobs or schedules 8-8
including unmanaged or non-Cisco in a VPN 25-11
inheriting policy rules 5-47
maps
adding existing managed 35-16
adding new managed 35-16
displaying devices from Device View 35-16
displaying managed 35-16
removing managed 35-16
showing containment for Catalyst switches, ASA, PIX, IPS devices 35-16
modifying policy assignment 5-49
modifying shared policies 5-49
naming conventions 3-3
overview of monitoring 1-7
policy status icons 5-30
preparing for management 2-1
property changes with policy effects 3-54
redeploying configuration files to 8-53

redeploying configurations to replaced hardware 8-53
renaming policies 5-48
replacing policies 5-44
rolling back configurations 8-69, 8-70, 8-71
selecting in site-to-site VPNs 25-32
selecting multiple 1-45
sharing multiple policies 5-42
sharing with PRSM 72-22
showing contained modules 3-56
system variables 7-7
testing connectivity 9-1
troubleshooting communication 9-8
troubleshooting communication and deployment 9-1
troubleshooting device discovery failures 3-7
unassigning policies 5-36
understanding out-of-band changes 8-12
unsharing policies 5-43
using global search to find specific devices 1-42
what counts as a device 3-3
device selector
filtering 1-45
Device Selector dialog box 1-45
Device Server Assignment dialog box 9-9
device status view
working with 3-64
Device Status View command 1-32
Device view
adding local rules to shared policies 5-45
assigning shared policies 5-44
cloning shared policies 5-47
configuring local policies 5-31
configuring VPN topologies 25-19
copying policies between devices 5-33
inheriting policies 5-47
managing policies 5-30
modifying policy assignments 5-49
modifying shared policies 5-49
overview 1-15
policy banner 5-38
policy shortcut menu 5-40
policy status icons 5-30
renaming policies 5-48
sharing local policies 5-41
sharing multiple policies 5-42
unassigning policies 5-36
understanding basic policy management 5-31
understanding shared policies 5-37
unsharing policies 5-43
device view
understanding 3-1
Device View command 1-32
Device Whitelist dialog box 19-15
DHCP
Cisco IOS routers
defining address pools 63-91
defining policies 63-90
DHCP Database dialog box 63-94
DHCP Policy page 63-92
IP Pool dialog box 63-94
overview 63-87
understanding database agents 63-88
understanding option 82 63-89
understanding relay agents 63-88
understanding secured ARP 63-89
configuring passthrough for IOS devices 23-3
PIX/ASA/FWSM 52-10
add/edit servers 52-12
advanced configuration 52-13
configuring DHCP servers 52-10
server options 52-13
traffic blocked 9-15
DHCP relay
interface-specific 46-41
Option 82 46-41, 52-5
PIX/ASA/FWSM 52-5, 52-7
add/edit agent 52-6
add/edit server 52-7

Trusted Interface (Option 82) 46-41, 52-5
DHCPv6 relay
PIX/ASA/FWSM
add/edit agent 52-9
add/edit server 52-9
diagnostics
setting debug options 11-11
diagnostics file, creating 10-28
dial backup
configuring in Easy VPN 28-2
configuring in VPN 25-40
configuring VPN advanced settings 25-41
Dial Backup Settings dialog box 25-41
dialer interfaces
defining BRI properties 62-30
defining profiles 62-28
Dialer Physical Interface dialog box 62-33
Dialer Policy page 62-31
Dialer Profile dialog box 62-32
on Cisco IOS routers 62-28
Diffie-Hellman groups
in IKE proposals 26-7
Digital Subscriber Line (DSL) 62-34
digital subscriber line-access multiplexer (DSLAM) 62-35
directed broadcasts
enabling 62-20
Disable/enable NAT rules 24-34, 24-46
Discard Activity command 1-36
Discard Activity dialog box 4-22
Discard command 1-31
Discard Deployment Job dialog box 8-20
Discard Ticket command 1-37
Discard Ticket dialog box 4-22
discovering
remote access VPNs 30-12
site-to-site VPNs 25-24
Discover Policies on Device command 1-32
Discover VPN Policies command 1-32
Discover VPN Policies wizard 25-24
discovery
default behavior settings 11-25
generating status report 10-30
invalid certificate error 9-7
overview 1-20
security certificate error 9-5, 9-6
setting debug options 11-11
Discovery Settings page 11-25
Discovery Status dialog box 5-23
discovery task
frequently asked questions 5-27
starting 5-15
viewing status 5-22
disk space, monitoring event data store 69-35
Display Actual Size command 1-33
Distributed Traffic Shaping (DTS) 66-7
DMVPN (Dynamic Multipoint VPN)
advantages of using with GRE 27-11
configuring 27-12
configuring GRE modes 27-12
large scale DMVPNs
configuring 27-16
configuring server load balancing 27-17
overview 27-1, 27-9
spoke-to-spoke connections 27-10
supported platforms 25-9
understanding 27-10
DNS
configuring for inspection rules 17-19
PIX/ASA/FWSM
add/edit server group 52-16
add server 52-17
servers page 52-14
DNS class map objects
creating 17-22
match criteria 17-36
DNS policy map objects
creating 17-22

match conditions and actions 17-36
properties 17-32
DNS servers
configuring for IPS global correlation 36-24
DNS snooping 19-6
dock
report windows 70-30
view windows 69-38
Dock Map View command 1-33
documentation
conventions 2-lxiii
Domain AD Server dialog box 13-10
Domain Name System (DNS)
Cisco IOS routers
defining policies 63-75
DNS Policy page 63-76
IP Host dialog box 63-76
overview 63-74
do not ask warnings, resetting 11-10
drill-down reports 70-26
DSLAM 62-35
duration
VPN user reports 70-16
dynamic access policies
attributes 32-4, 32-7
configuring 32-2
managing 32-1
understanding 32-1
dynamic access policies (DAP) 32-30
Dynamic Access Policy page
Add/Edit Dynamic Access Policy dialog box
Add/Edit DAP Entry dialog box 32-21
Add/Edit DAP Entry dialog box > AAA Attributes Cisco 32-22
Add/Edit DAP Entry dialog box > AAA Attributes LDAP 32-24
Add/Edit DAP Entry dialog box > AAA Attributes RADIUS 32-25
Add/Edit DAP Entry dialog box > Anti-Spyware 32-26
Add/Edit DAP Entry dialog box > Anti-Virus 32-27
Add/Edit DAP Entry dialog box > AnyConnect Identity 32-28
Add/Edit DAP Entry dialog box > Application 32-29
Add/Edit DAP Entry dialog box > File 32-31
Add/Edit DAP Entry dialog box > NAC 32-32
Add/Edit DAP Entry dialog box > Operating System 32-33
Add/Edit DAP Entry dialog box > Personal Firewall 32-34
Add/Edit DAP Entry dialog box > Policy 32-35
Add/Edit DAP Entry dialog box > Process 32-36
Add/Edit DAP Entry dialog box > Registry 32-37
Advanced Expressions tab 32-44
Logical Operations tab 32-42
Main tab 32-14
Dynamic Access Policy page (ASA) 32-11
Cisco Secure Desktop Manager Policy Editor dialog box 32-46
Dynamic Access policy page (ASA) > Add/Edit Dynamic Access Policy dialog box 32-12
Dynamic Blacklist Configuration tab 19-10
dynamic crypto maps 26-19
dynamic filter snooping (DNS)
enabling 17-19
Dynamic Multipoint VPN (DMVPN)
mandatory and optional policies 25-6
dynamic NAT
Cisco IOS routers 24-10
Dynamic Translation Rule
PIX/ASA/FWSM 24-22
add/edit 24-23
dynamic VTI
configuring in Easy VPN 28-12
in remote access VPNs 33-7
understanding use in Easy VPN 28-3

E

Easy VPN
- configuration modes 28-3
- configuration overview 28-5
- configuring client connection characteristics 28-7
- configuring dial backup 28-2
- configuring dynamic VTI 28-12
- configuring high availability 28-2
- connection profile policies 28-13
- connection profiles (ASA, PIX 7+) 31-8
- extended authentication (xauth) 28-4
- important configuration notes 28-6
- IPsec proposals 28-10
- mandatory and optional policies 25-6
- overview 28-1
- supported platforms 25-9
- understanding 28-1
- understanding dynamic VTI 28-3
- user group policies 28-14

ECMP 22-4
Edit AAA Option dialog box 15-19
Edit AAA Rule dialog box 15-13
Edit AAA Server dialog box 6-33
Edit AAA Server Group dialog box 6-49
Edit Access Rule dialog box 16-14
Edit Actions dialog box 39-12
Edit activity state 4-4
Edit AOL Class Map dialog box 17-28, 21-19
Edit A Port Forwarding Entry dialog box 34-41
Edit ASA Group Policies dialog box
- client configuration settings 34-6
- client firewall attributes 34-7
- connection settings 34-33
- DNS/WINS settings 34-29, 34-30
- hardware client attributes 34-9
- IPSec settings 34-10
- overview 34-1
- split tunneling settings 34-31
- SSL VPN clientless settings 34-12
- SSL VPN full client settings 34-19
- SSL VPN settings 34-25
- technology settings 34-1

Edit A Smart Tunnel Entry dialog box 34-67, 34-70
Edit AS Path Entry dialog box 56-152
Edit AS Path Object dialog box 56-151
Edit Auto Signon Rules dialog box 34-27
Edit Auto Update Settings dialog box 11-54
Edit Category dialog box 12-14
Edit Cisco Secure Desktop Configuration dialog box 34-35
Edit Client Access Rules dialog box 34-12
Edit Client Update dialog box 34-81
Edit Column dialog box 34-61
Edit Community List Entry dialog box 56-154, 56-155
Edit Community List Object dialog box 56-153
Edit Custom Pane dialog box 34-62
Edit DCE/RPC Map dialog box 17-29
Edit Deploy Method dialog box 8-30
Edit Description dialog box 12-14
Edit Destinations dialog box 12-11
Edit Device Groups command 1-31
Edit Device Groups dialog box 3-61
Edit DNS Class Map dialog box 17-28
Edit DNS Map dialog box
- Filtering tab 17-34, 17-35
- overview 17-32
- Protocol Conformance tab 17-33

Edit eDonkey Class Map dialog box 17-28, 21-19
Edit Endpoints dialog box
- FWSM tab 25-47
- overview 25-34
- Protected Networks tab 25-46
- VPN Interface tab 25-36, 25-50
- VPNSM/VPN SPA/VSPA settings, VPN Interface tab 25-42
- VRF Aware IPsec tab 25-48

Edit ESMTP Map dialog box 17-39

Edit Extended Access Control Entry dialog box 6-61
Edit Extended Access List dialog box 6-59
Edit External Filter dialog box 21-41
Edit Extranet VPN dialog box
overview 25-66
Edit FastTrack Class Map dialog box 17-28, 21-19
Edit Fidelity dialog box 39-13
Edit File Object dialog box 34-37
Edit FlexConfig dialog box 7-30
Edit FTP Class Map dialog box 17-28
Edit FTP Map dialog box 17-42
Edit Gnutella Class Map dialog box 17-28, 21-19
Edit Group Member dialog box 29-21
Edit GTP Map dialog box 17-45
Edit H.323 Class Map dialog box 17-28, 21-19
Edit H.323 Map dialog box 17-51, 21-34
Edit HSI Endpoint IP Address dialog box 17-54
Edit HSI Group dialog box 17-53
Edit HTTP Class Map dialog box 17-28, 21-19
Edit HTTP Map dialog box 21-34
ASA 7.1.x, PIX 7.1.x, FWSM 3.x, IOS devices
Entity Length tab 17-58
Extension Request Method tab 17-61
General tab 17-57
overview 17-56
Port Misuse tab 17-62
RFC Request Method tab 17-60
Transfer Encoding tab 17-63
ASA 7.2+ and PIX 7.2+ devices 17-64
Edit ICQ Class Map dialog box 17-28, 21-19
Edit IKEv1 Proposal dialog box 26-10
Edit IKEv2 Proposal dialog box 26-14
Edit IMAP Class Map dialog box 17-28, 21-19
Edit IMAP Map dialog box 21-34
Edit IM Class Map dialog box 17-28
Edit IM Map dialog box 21-34
ASA and PIX device 17-70
IOS device 17-73
Edit Inspect/Application FW Rule wizard
Address and Port page 17-13
Inspected Protocol page 17-17
Match Traffic page 17-11
Edit Inspect Parameter Map dialog box 21-31
Edit Interfaces dialog box 12-13
Edit IP Options Map dialog box 17-75
Edit IPsec Pass Through Map dialog box 17-80
Edit IPSec Transform Set dialog box 26-27
Edit IPv4 Pool Object dialog box 6-92
Edit IPv6 Map dialog box 17-77, 17-91
Edit IPv6 Pool Object dialog box 6-93
Edit Kazaa2 Class Map dialog box 17-28, 21-19
Edit Key Server dialog box 29-19
Edit Language dialog box 34-56
Edit LDAP Attribute Map dialog box 6-46
Edit LDAP Attribute Map Value dialog box 6-47
Edit Load Balancing Parameters dialog box 27-17
Edit Local Web Filter Class Map dialog box 17-28, 21-19
Edit Local Web Filter Parameter Map dialog box 21-38
Edit MAC Address Pool Object dialog box 6-94
Edit Map Value dialog box 6-47
Edit Match Condition and Action dialog box
DNS policy maps 17-36
ESMTP policy maps 17-40
FTP policy maps 17-43
GTP policy maps 17-49
H.323 (IOS) policy maps 21-35
H.323 policy maps 17-54
HTTP (Zone Based IOS) policy maps 21-35
HTTP policy maps 17-66
IM (Zone Based IOS) policy maps 21-35
IMAP policy maps 21-35
IM policy maps 17-71
IPv6 policy maps 17-78, 17-92
P2P policy maps 21-35
POP3 policy maps 21-35
SIP (IOS) policy maps 21-35
SIP policy maps 17-85, 17-95
Skinny policy maps 17-89

SMTP policy maps 21-35
Sun RPC policy maps 21-35
Web Filter policy maps 21-35
Edit Match Criterion dialog box
AOL class maps 21-21
DNS class maps 17-36
eDonkey class maps 21-21
FastTrack class maps 21-21
FTP class maps 17-43
Gnutella class maps 21-21
H.323 (IOS) class maps 21-22
H.323 class maps 17-54
HTTP (IOS) class maps 21-22
HTTP class maps 17-66
ICQ class maps 21-21
IMAP class maps 21-25
IM class maps 17-71
Kazaa2 class maps 21-21
Local Web Filter class maps 21-29
MSN Messenger class maps 21-21
N2H2 class maps 21-30
POP3 class maps 21-25
SIP (IOS) class maps 21-25
SIP class maps 17-85, 17-95
SMTP class maps 21-27
Sun RPC class maps 21-29
Websense class maps 21-30
Windows Messenger class maps 21-21
Yahoo Messenger class maps 21-21
Edit menu
Configuration Manager 1-31
Edit MSN Messenger Class Map dialog box 17-28, 21-19
Edit N2H2 Parameter Map dialog box 21-39
Edit N2H2 Web Filter Class Map dialog box 17-28, 21-19
Edit NAT Rule dialog box
ASA 8.3+ 24-36
Edit NetBIOS Map dialog box 17-81
Edit Network/Host dialog box
General tab 6-83
NAT tab 24-42
Edit Options dialog box 16-17
Edit P2P Map dialog box 21-34
Edit Permit Response dialog box 17-48
Edit Per-Session NAT Rule dialog box 24-47
Edit PIX/ASA/FWSM Web Filter Rule dialog box 18-5
Edit PKI Enrollment dialog box
CA Information tab 26-60
Certificate Subject Name tab 26-66
Enrollment Parameters tab 26-63
overview 26-58
Trusted CA Hierarchy tab 26-67
Edit Policy Assignments command 1-32
Edit Policy List Object dialog box 56-143
Edit POP3 Class Map dialog box 17-28, 21-19
Edit Port Forwarding List dialog box 34-40
Edit Port List dialog box 6-102
Edit Prefix List Entry dialog box 56-148, 56-150
Edit Prefix List Object dialog box 56-146, 56-148
Edit Protocol Info Parameter Map dialog box 21-33
Edit Regular Expression dialog box 17-108
Edit Regular Expression Group dialog box 17-108
Edit Route Map Entry dialog box 56-137
Edit Route Map Object dialog box 56-136
Edit Row command 1-31
Edit Rule Section dialog box 12-22
Edit Security Association Dialog Box 25-58
Edit Selected Deployment Method dialog box 8-30
Edit Server dialog box
Protocol Info Parameter maps 21-34
Edit Server Group dialog box 15-19
Edit Service dialog box 6-103
Edit Services dialog box 12-13
Edit Signature dialog box 39-15
Edit Signature Parameter—Component List dialog box 39-29
Edit Signature Parameters dialog box 39-24
Edit Single Sign On Server dialog boxes 34-42
Edit SIP Class Map dialog box 17-28, 21-19

Edit SIP Map dialog box 17-83, 17-93, 17-102, 17-103, 21-34
Edit Skinny Map dialog boxes 17-87
Edit SLA Monitor dialog box 51-10
Edit Smart Tunnel Auto Signon Entry dialog box 34-72
Edit Smart Tunnel Auto Signon Lists dialog box 34-71
Edit Smart Tunnel Lists dialog box 34-66, 34-69
Edit SMTP Class Map dialog box 17-28, 21-19
Edit SMTP Map dialog box 21-34
Edit SNMP Map dialog box 17-90
Edit Sources dialog box 12-11
Edit SSL VPN Customization dialog box 34-51
  Applications 34-60
  Copyright Panel 34-58
  Custom Panes 34-61
  Full Customization 34-59
  Home Page 34-62
  Informational Panel 34-57
  Language 34-54
  Logon Form 34-56
  Logout Page 34-63
  Title Panel 34-53
  Toolbar 34-59
Edit SSL VPN Gateway dialog box 34-64
Edit Standard Access Control Entry dialog box 6-64
Edit Standard Access List dialog box 6-59
Edit Sun RPC Class Map dialog box 17-28, 21-19
Edit Sun RPC Map dialog box 21-34
Edit TCP Map dialog box 58-22
Edit TCP Option Range Dialog Box 58-25
Edit Text Object dialog box 7-32
Edit Time Range dialog box 6-71
Edit Traffic Flow dialog box 58-18
Edit Translated Address dialog box 24-29
Edit Transparent EtherType dialog box 23-7
Edit Transparent Firewall Rule dialog box 23-5
Edit Transparent Mask dialog box 23-7
Edit Trend Content Filter Class Map dialog box 17-28, 21-19
Edit Trend Parameter Map dialog box 21-42
Edit Unified Access Control Entry dialog box 6-67
Edit Update Server Settings dialog box 11-52
Edit URL Domain Name dialog box 21-45
Edit URLF Glob Parameter Map dialog box 21-45
Edit URL Filter Parameter Map dialog box 21-43
Edit User Credentials dialog box 36-19
Edit User dialog box 12-12
Edit User Group dialog box
  Advanced PIX 6.3 settings 34-82
  Browser Proxy settings 34-87
  Client (IOS) settings 34-78
  Clientless settings 34-83
  Client VPN Software Update (IOS) settings 34-81
  DNS/WINS settings 34-77
  General settings 34-75
  IOS Xauth Options settings 34-80
  overview 34-73
  Split Tunneling settings (Easy VPN/remote access IPSec VPN) 34-77
  SSL VPN Connection settings 34-88
  SSL VPN Full Tunnel settings 34-84
  SSL VPN Split Tunneling settings 34-86
  Technology settings 34-73
  Thin Client settings 34-84
Edit VDI Server dialog box 34-15
Edit Virtual Sensor dialog box 38-7, 38-8
Edit VPN dialog box
  Device Selection tab 25-32
  Edit Endpoints dialog box 25-34
  Endpoints tab 25-34
  High Availability tab 25-52
  Name and Technology tab 25-30
  overview 25-28
Edit Web Access Control Entry dialog box 6-65
Edit Web Filter Map dialog box 21-47
Edit Web Filter Options dialog box 18-9
Edit Web Filter Type dialog box 18-8
Edit Websense Parameter Map dialog box 21-39

Edit Websense Web Filter Class Map dialog box 17-28, 21-19
Edit Web Type Access List dialog box 6-59
Edit Windows Messenger Class Map dialog box 17-28, 21-19
Edit WINS Server dialog box 34-90
Edit WINS Server List dialog box 34-89
Edit Yahoo Messenger Class Map dialog box 17-28, 21-19
Edit Zones dialog box 12-13
eDonkey class map objects
creating 21-16
match criteria 21-21
EIGRP routing
defining interface properties 67-10
defining routes 67-9
EIGRP Routing Policy page 67-13
Interface dialog box 67-16
Interfaces tab 67-15
on Cisco IOS routers 67-8
PIX/ASA/FWSM
advanced settings 56-34
Filter Rule configuration 56-40
Filter Rules tab 56-39
Interface configuration 56-48
Interfaces tab 56-47
neighbor configuration 56-42
Neighbors tab 56-41
policy 56-32
redistribution configuration 56-44
Redistribution tab 56-42
Setup tab 56-36
Summary Address configuration 56-46
Summary Address tab 56-45
redistributing routes 67-12
Redistribution Mapping dialog box 67-18
Redistribution tab 67-17
Setup dialog box 67-14
Setup tab 67-13
e-mail
blocking spam using zone-based firewall rules 21-27
preventing DoS attacks 21-27
e-mail notifications
configuring SMTP server 1-27
PIX/ASA/FWSM
recipient set-up 54-8
syslog messages 54-8
embedded event manager
add/edit action configuration 54-7
add/edit applet 54-5
add/edit syslog configuration 54-7
ASA 54-3
Enable/disable NAT rules 24-34, 24-46
Enable PIM and IGMP
PIX/ASA/FWSM 55-1
Encapsulating Security Protocol (ESP) encryption algorithm 26-29
encoding rules
defining for SSL VPN (ASA) 31-55
encryption algorithms
3DES (Triple DES) 26-6
AES (Advanced Encryption Standard) 26-6
DES (Data Encryption Standard) 26-6
in IKE proposals 26-6
endpoints and protected networks
configuring dial backup 25-40
defining in GET VPN topologies 25-60
defining in VPN topologies 25-34
VPN Interface tab 25-36, 25-50
equal-cost multi-path 22-4
Error Writing to Server deployment errors 9-15
ESMTP
configuring for inspection rules 17-20
ESMTP policy map objects
creating 17-22
match conditions and actions 17-40
properties 17-39
EtherChannel

Create and Edit IDSM EtherChannel VLANs dialog boxes 68-49
defining IDSM VLANs 68-44
deleting IDSM VLANs 68-45
EtherChannels
ASA 46-9
edit assigned interface 46-12
LACP 46-12
load balancing 46-13
evaluation license
upgrading to permanent license 10-16
event
lists 54-9
add/edit 54-10
syslog class
add/edit 54-11
syslog message ID
add/edit 54-11
Event Action Filters page 40-7
Event Action Overrides page 40-13
event actions, IPS
configuring filter rules 40-4
configuring network information 40-17
configuring OS maps 40-21
configuring overrides 40-13
configuring settings 40-23
configuring target value ratings 40-17
example filter rule 69-67
filter rule attributes 40-9
filter rules policy 40-7
filter rules tips 40-6
overview 40-1
possible actions 40-2
process overview 40-1
Event Management page 11-27, 11-35
CPU Throttling Policy dialog box 11-33
event manager applet 54-3
Event Manager service
configuring 69-30
managing 69-30
monitoring event store disk space 69-35
monitoring status 69-31
selecting devices to monitor 69-34
starting and stopping 69-30
status icon colors 69-31
events
archiving (backing up) the event data store 69-36
configuring firewall devices (ASA, FWSM) 69-28
configuring IPS devices 69-29
copying 69-53
CS-MARS 72-46
looking up 72-41
looking up policies based on related events 72-45
Netflow support for policy lookup 72-47
viewing access rule events 72-42
viewing IPS signature events 72-44
ensuring time synchronization 69-27
Event Viewer
clearing filters 69-48
context menu 69-49
cross-launching from HPM 69-58
filtering by column 69-45
filtering by events 69-47
filtering overview 69-43
looking up 69-55
looking up policies based on related events 69-54
refreshing event table 69-44
selecting time range 69-43
text searches (quick filter) 69-47
using time slider with filtering 69-44
viewing access rule events 69-56
viewing IPS signature events 69-57
examining details 69-53
examples of analysis
mitigating botnet activity 69-65
monitoring and mitigating botnet activity 69-61
monitoring botnet activity using ASDM 69-64

monitoring botnet activity using Event Viewer 69-62
monitoring botnet activity using Report Manager 69-64
monitoring identity-aware firewall policies 13-27
monitoring TrustSec policies 14-17
overview 69-58
removing false positive IPS events 69-66
understanding botnet syslog events 69-61
user access to server blocked 69-59
performing operations on 69-49
properties 69-18
recovering the event data store 69-36
saving to a file 69-53
understanding Event Viewer access control 69-4
viewing 69-1
Event Viewer
archiving (backing up) the event data store 69-36
arranging views 69-38
ASA devices, configuring to provide events 69-28
columns 69-18
configuring color rules 69-40
configuring Event Manager service 69-30
copying events 69-53
creating custom views 69-41
cross-launching from HPM 69-58
deleting custom views 69-43
editing view name and description 69-41
ensuring time synchronization 69-27
Event Monitoring window 69-14
events
context menu 69-49
historical and real-time lookup 69-55
looking up 69-55
event table
customizing appearance 69-39
event details pane 69-26
refreshing 69-44
time slider 69-25
toolbar 69-16
examining event details 69-53
examples of analysis
mitigating botnet activity 69-65
monitoring and mitigating botnet activity 69-61
monitoring botnet activity 69-62
monitoring identity-aware firewall policies 13-27
monitoring TrustSec policies 14-17
overview 69-58
removing false positive IPS events 69-66
understanding botnet syslog events 69-61
user access to server blocked 69-59
features
historical views 69-2
overview 69-1
policy navigation 69-3
real-time views 69-2
views and filters 69-3
File menu reference 69-9
filters
advantages of using network/host objects 69-67
clearing 69-48
column based 69-45
event based 69-47
overview 69-43
submission requirements for policy objects 69-68
text searches (quick filter) 69-47
time range 69-43
time slider 69-44
floating views 69-38
FWSM devices, configuring to provide events 69-28
IPS devices, configuring to provide events 69-29
limits of 69-4
looking up Security Manager policies based on events 69-54
managing service 69-30
monitoring event store disk space 69-35
monitoring status 69-31
opening views 69-38

overview 69-7
performing operations on 69-49
preparation for use 69-27
recovering the event data store 69-36
saving events 69-53
saving views 69-42
selecting devices to monitor 69-34
settings 11-27, 11-35
starting or stopping the Event Manager service 69-30
status icon colors 69-31
switching between IP addresses and host object names 69-39
switching between real-time and historical views 69-42
syslogs 69-6
troubleshooting
Event Viewer Unavailable message 11-27, 11-36, 69-30
policy objects not available for filtering 69-68
understanding access control 69-4
using 69-37
using views 69-37
viewing access rule events 69-56
viewing IPS signature events 69-57
view list 69-12
View menu reference 69-10
Event Viewer command 1-38
exclusive domains
configuring for IOS devices 18-10
Exit command 1-31
Exit command (Report Manager) 70-8
exiting
Cisco Security Management Suite server 1-12
CiscoWorks Common Services 1-12
Security Manager 1-11, 1-12
expiration dates
configuring for access rules 16-22
export
device inventory
DCR, CS-MARS, Security Manager formats 10-6
device with policies 10-6
overview 10-6
supported CSV formats 10-9
HPM data 71-31
IPS event action overrides 40-13
IPS event filter rules 40-4, 40-7
policy objects 6-23
reports 70-28
shared policies 10-12
Export Devices or Policies commands 1-30
Export Inventory dialog box 10-6
Export Map command 1-33
External Product Interface dialog box 36-27
External Product Interface policy 36-26

F

factory-default configurations 46-2
failover
Active/Active
command replication 50-4
configuration synchronization 50-3
add new context to group 2 50-8
configuring in site-to-site VPN 25-52
edit bridge group 50-17
FWSM 50-13
advanced settings 50-16
PIX/ASA 50-17
Add Failover Group 50-25
settings 50-21
PIX/ASA/FWSM 50-10
active/active 50-2, 50-3
active/standby 50-2
bootstrap configuration 50-26
configuration basics 50-5
configuring 50-1
interface configuration 50-23

interface MAC address 50-23
security context 50-26
stateful 50-3, 50-4
stateless 50-3
types of 50-2
understanding 50-1
PIX 6.3 50-10
interface configuration 50-12
stateful in site-to-site VPN 25-54
false negatives
definition of 39-23
false positives
definition of 39-23
FastTrack class map objects
creating 21-16
match criteria 21-21
feature sets 1-4
File menu
Configuration Manager 1-30
Event Viewer 69-9
Report Manager 70-8
file objects
attributes 34-37
selecting 34-39
files
deploying to 8-11
selecting or specifying 1-50
Filter Item dialog box 40-9
filter rules, event action (IPS)
attributes 40-9
configuring 40-4
example rule 69-67
exporting 40-4
policy 40-7
tips 40-6
filters
Event Viewer
clearing 69-48
column based 69-45
context menu 69-49
event based 69-47
overview 69-43
refreshing event list 69-44
selecting time range 69-43
text searches (quick filter) 69-47
using time slider 69-44
filtering selectors 1-45
filtering tables 1-48
HPM
column based 71-17
custom 71-18
filters (Event Viewer)
advantages of using network/host objects 69-67
overview 69-3
submission requirements for policy objects 69-68
Find and Replace dialog box 12-17
find and replace in rules policies 12-16
Find Map Node command 1-33
Find Node dialog box 35-12
FirePOWER
ASA module
detecting 72-21
FireSIGHT Management Center
starting from Security Manager 72-20
FireSIGHT Management Center command 1-37
Firewall
AAA IOS Timeout Values 15-30
firewall
AAA firewall
advanced settings 15-20
configuring 15-6
MAC exempt lists 15-26
AAA firewall policy
advanced settings 15-20
configuring 15-6
AAA page 15-28
AAA rules
configuring AAA firewall settings 15-6

configuring AuthProxy settings 15-9
configuring cut-through proxy (ASA) 13-23
configuring for ASA/PIX/FWSM devices 15-4
configuring for IOS devices 15-7
configuring identity aware 13-21
configuring security group aware 14-17
managing 15-1
properties 15-13
understanding 15-1
understanding how users authenticate 15-2
Access Control page 16-24
access controls
per user downloadable ACLs 16-27
access control settings
configuring settings 16-23
access rule
event analysis example, user access blocked 69-59
finding from CS-MARS events 72-45
finding from Event Viewer events 69-54
viewing related CS-MARS events 72-42
viewing related events 69-56
access rules
address requirements 16-5
configuring 16-7
configuring expiration dates 16-22
configuring identity aware 13-21
configuring security group aware 14-17
how deployed 16-5
import examples 16-44
importing 16-40
IPS blocking, affect of 43-4
managing 16-1
optimizing during deployment 16-46
sharing ACLs among interfaces 11-18
understanding 16-1
understanding device-specific behavior 16-4
understanding global 16-3
understanding requirements when using inspection 17-4
ACL naming conventions 12-5
adding rules 12-9
analysis reports 16-34
AuthProxy
configuring 15-9
AuthProxy settings policy
configuring 15-9
botnet traffic filter rules 19-9
combining rules
example 12-27
interpreting results 12-25
procedure 12-22
configuring policies in Map view 35-23
configuring settings 18-15
configuring settings policies in Map view 35-23
conflict detection 16-28
converting IPv4 rules 12-28
deleting rules 12-9
device types 46-1
disabling rules 12-20
editing rules 12-10
enabling rules 12-20
finding and replacing items in rules policies 12-16
Firewall ACL Setting dialog box 16-26
identity-aware policies
collecting user statistics 13-25
configuring 13-7
configuring cut-through proxy 13-23
configuring identity options 13-15
configuring rules 13-21
configuring the ASA 13-7, 14-8
enabling 13-8
filtering VPN traffic 13-26
identifying AD servers and agents 11-38, 13-8
managing 13-1
monitoring 13-27
overview 13-1

requirements 13-3
user identity acquisition 13-2
Inspection page 17-111
inspection rules
add/edit rule wizard 17-11, 17-13, 17-17
choosing interfaces 17-2
configuring 17-5
configuring identity aware 13-21
configuring security group aware 14-17
managing 17-1
preventing DoS attacks on IOS devices 17-5
selecting protocols 17-3, 17-17
understanding 17-2
understanding access rule requirements 17-4
inspection settings
configuring for IOS devices 17-111
introduction 12-1
IPv6 access rules
configuring expiration dates 16-22
sharing ACLs among interfaces 11-18
understanding global 16-3
MAC exempt lists, AAA firewall 15-26
managing rules tables 12-7
moving rules 12-19
object groups
expanding during discovery 12-35
optimizing network object groups during deployment 12-35
overview 12-1
per user downloadable ACLs 16-27
policy discovery 5-13
policy query
example report 12-34
generating reports 12-28
interpreting results 12-32
preserving ACL names 12-4
reference information for AAA rules 15-20
resolving access rule conflicts 16-34
resolving ACL naming conflicts 12-7
rule table sections 12-20
security group aware policies
configuring ISE settings 11-56
configuring rules 14-17
security group-aware policies
configuring 14-7
managing 14-1
system variables 7-9
transparent rules
adding or editing a rule 23-5
configuring 23-1
configuring passthrough for IOS devices 23-3
editing the EtherType 23-7
editing the mask 23-7
managing 23-1
Transparent Rules page 23-3
TrustSec firewall policies
configuring 14-7
managing 14-1
overview 14-1
TrustSec policies
monitoring 14-17
understanding NAT effects 12-3
understanding rule order 12-19
understanding rule processing order 12-2
using rules tables 12-8
Web Filter page 18-16
web filter rules
configuring for ASA, PIX, FWSM devices 18-2
configuring for IOS devices 18-10
managing 18-1
understanding 18-1
zone-based firewall
add/edit zones 21-53
advanced options 21-67
configuring PAM 21-69
configuring rules 21-13, 21-62
configuring settings 21-49
Content Filter tab 21-52

designing network zones 21-1
development overview 21-12
Global Parameters tab 21-50
page 21-50
protocol selection 21-68
rules table 21-58
tabs 21-49
VPN tab 21-50
WAAS tab 21-50
Zones tab 21-50
zone-based firewalls
changing the default drop rule 21-48
general recommendations 21-12
IPSec VPN 21-6
logging 21-1
overview 21-1
restrictions 21-3
Self zone 21-5
troubleshooting 21-54
understanding 21-3
understanding permit/deny and action 21-8
understanding services and protocols 21-11
VRF 21-7
Firewall AAA IOS Timeout Value Setting dialog box 15-30
Firewall AAA MAC Exempt Setting dialog box 15-27
Firewall ACL Setting dialog box 16-26
Firewall Device dialog box 43-14
Firewall Services Module
see FWSM 47-1
Fit to Window command 1-33
FlexConfig objects
adding to policies 7-35
ASA samples 7-20
Catalyst 6500/7600 samples 7-22
changing order in policies 7-35
changing variable values 7-35
Cisco IOS Software samples 7-22
CLI commands 7-2
configuring 7-25
configuring AAA for administrative introducers 63-84
creating 7-28
creating text objects 7-32
deleting variables 7-28
PIX firewall samples 7-23
previewing CLI 7-35
properties 7-30
property selector 7-34
removing from policies 7-35
router samples 7-24
samples 7-19
scripting language
example of looping 7-3
example of looping with if/else statements 7-4
example of two-dimensional looping 7-3
understanding 7-3
system variables
device 7-7
firewalls 7-9
remote access VPN 7-19
router 7-13
understanding 7-7
VPN 7-14
undefined variables 7-33
understanding 7-2
variables 7-5
variables, example 7-6
FlexConfig policies
adding objects 7-35
changing object order 7-35
changing variable values 7-35
configuring 7-25
configuring AAA for administrative introducers 63-84
editing 7-35
previewing CLI 7-35
removing objects 7-35

understanding 7-2
FlexConfig Policy page 7-36
FlexConfig Preview dialog box 7-38
FlexConfigs
creating (scenario) 7-25
managing 7-1
troubleshooting 7-38
FlexConfig Undefined Variables dialog box 7-33
float
report windows 70-30
view windows 69-38
floodguard 57-2
FQDN objects
creating 6-82
understanding 6-80
fragmentation
configuring settings in VPNs 26-31, 26-44
fragments settings 57-2
frequently asked questions
policy discovery 5-27
FTP class map objects
creating 17-22
match criteria 17-43
FTP policy map objects
creating 17-22
match conditions and actions 17-43
properties 17-42
full mesh topologies
description 25-4
partial mesh 25-5
full tunnel client access mode 30-5
FWSM
AAA support 6-28
about 46-1
adding SSL thumbprints manually 9-5
adding when using multiple-context mode 3-7
adding when using non-default HTTPS (SSL) port 3-7
Asymmetric Routing Groups 46-6
Bridge Groups
add/edit 46-62
bridge groups 47-3
changing deployment method to serial for multiple-context mode 9-17
configuring for event management 69-28
configuring FWSM endpoints in site-to-site VPNs 25-47
configuring transparent firewall rules 23-1
credentials 3-19
deleting security contexts 59-7
deployment failures after changing interface policies 9-16
deployment failures in multiple-context mode 9-16
deployment failures with large ACLs 9-17
Device Access
managing Resources 51-2
Resources 51-3
Resources, add/edit 51-4
discovering failover modules 3-7
Event Viewer support 69-4
Failover 50-13
advanced settings 50-16
edit bridge group 50-17
including in deployment jobs 8-28
interfaces
add/edit 46-31
configuring 46-3
General tab 46-33
IPv6 46-47, 46-73
IPv6, add/edit 46-52
IPv6, add/edit prefixes 46-54, 46-56
managing 46-26
packet capture, using 72-30
PDM 72-15
policy discovery 5-13
rollback, commands to recover from failover misconfiguration 8-68
rollback command conflicts 8-67
rollback restrictions for failover devices 8-65

rollback restrictions for multiple context mode 8-64
security contexts
configuration 59-8
selecting policy types to manage 5-11
setting up SSL (HTTPS) 2-3
SSL certificate configuration 11-22
TCP State Bypass 58-3
troubleshooting deployment 9-16

G

General
PIX/ASA/FWSM
security policies 57-1
General Configuration tab, SNMP policy for IPS 36-10
General page, device properties 3-41
General tab, IPS blocking policy 43-10
General tab (Translation Rules)
PIX/ASA/FWSM 24-31
generic routers 3-8
GET VPN
anti-replay, time based 29-11
configuring 29-12
configuring global ISAKMP and IPsec settings 29-16
configuring group members 29-20
cooperative key servers 29-7
defining group encryption 25-54
generating, synchronizing RSA keys 29-13
group members
adding 29-19
editing 29-21
IKE proposal 29-15
key servers
adding 29-19
editing 29-19
mandatory and optional policies 25-6
migrating to 29-23
overview 29-1
receive-only SAs 29-23
registration
choosing the rekey transport mechanism 29-6
configuring fail-close mode 29-8
registration process 29-4
SAs
passive SA mode 29-23
receive-only mode 29-23
security policy 29-10
supported platforms 25-9
troubleshooting 29-25
understanding 29-2
GET VPNs
group encryption policies
certificate authorization 25-58
security associations 25-58
global correlation
configuring 42-1
configuring DNS servers 36-24
configuring HTTP proxy server 36-24
configuring inspection and reputation 42-5
configuring network participation 42-7
configuring with Botnet Traffic Filtering 42-1
data collected 42-3
requirements and limitations 42-4
understanding 42-1
understanding network participation 42-3
understanding reputation 42-2
Global Search
using 1-42
Global Search command 1-31
global settings
remote access VPN
configuring 26-30
Gnutella class map objects
creating 21-16
match criteria 21-21
GRE (generic routing encapsulation) VPN
advantages of IPsec tunneling with GRE 27-3
configuring 27-5

configuring GRE modes 27-6
dynamically addressed spokes 27-5
implementation 27-3
overview 27-1, 27-2
prerequisites for successful configuration 27-3
supported platforms 25-9
understanding 27-2
GRE Dynamic IP
mandatory and optional policies 25-6
GRE Modes Page
DMVPN properties 27-12
GRE or GRE Dynamic IP properties 27-6
overview 27-1
Group Domain of Interpretation (GDOI) protocol 29-3
group encryption
defining in GET VPN topologies 25-54
Group Encryption Policy page (GET VPN) 25-54
group members
adding 29-19
communication flow 29-2
configuring fail-close mode 29-8
editing 29-21
GET VPN
registration process 29-4
security policy ACLs 29-10
group members (GET VPN)
configuring 29-20
Group Members page (GET VPN) 29-20
group policies
configuring 31-26
creating 31-28
understanding 31-27
VPNs
configuring bookmarks 31-82
configuring portal appearance 31-78
configuring WINS servers for file system access 31-88
customizing 31-77
post URL method and macro substitutions in bookmarks 31-84
smart tunnels 31-85
Group Policies page 31-26
groups
adding or removing devices 3-63
creating 3-62
deleting 3-63
understanding 3-60
working with 3-59
group types
creating 3-62
deleting 3-63
GTP map objects
Add Country Network Codes dialog box 17-48
Edit Country Network Codes dialog box 17-48
GTP Map Timeouts dialog box 17-48
GTP policy map objects
creating 17-22
match conditions and actions 17-49
properties 17-45

H

H.323 class map objects
IOS
creating 21-16
match criteria 21-22
match criteria 17-54
H.323 policy map objects
ASA/PIX/FWSM
creating 17-22
properties 17-51
IOS
creating 21-16
match conditions and actions 21-35
match conditions and actions 17-54
hash algorithms
in IKE proposals 26-6

MD5 26-7
SHA 26-6
Health & Performance Monitor command 1-38
Health and Performance Monitor
see HPM 71-1
viewing related events in Event Viewer 69-58
Health and Performance Monitor in Dashboard 72-2
help
accessing 1-52
Help About This Page command 1-38
helper addresses 62-14
Help menu
Configuration Manager 1-38
Help Topics command 1-38
Hide Navigation Window command 1-33
high availability (HA groups)
configuring in Easy VPN 28-2
configuring in site-to-site VPN 25-52
stateful/stateless failover 25-54
high availability policies
configuring in remote access VPNs 33-11
Histogram dialog box 41-13
histograms
configuring anomaly detection 41-11
understanding anomaly detection 41-9
Hit Count Details
example 16-38
Hit Count Details page 16-36
Hit Count Selection Summary Dialog Box 16-20
Hostname
PIX/ASA/FWSM 51-1
hostnames
Cisco IOS routers
defining 63-77
Hostname Policy page 63-78
overview 63-77
HPM
access control 71-3
Alerts
firewall 71-37
IPS 71-35
VPN 71-39
VPN, SNMP configuration 71-40
alerts 71-32
acknowledging 71-42
clearing 71-42
configuring 71-34
history 71-43
viewing 71-41
application window 71-6
Alerts display 71-32
Monitoring display 71-25
columns
Alert table 71-16
Device-related 71-8
showing/hiding 71-8
sorting 71-8
VPN-related 71-13
configuring for 71-4
custom views 71-24
device
monitoring 71-21
monitoring multiple contexts 71-3
priority monitoring 71-32
views 71-21
Device Manager
launching 71-3, 71-27
device manager
cross-launch 71-32
devices
managing 71-5
email notifications
configuring 71-34
export data 71-31
filters
column based 71-17
introduction 71-1
launching 71-4

List Filter 71-19
monitoring
device details 71-28
device status list 71-27
RA and S2S views 71-30
Summary 71-27
VPN details 71-28
VPN Summary list 71-27
overview 71-1
read time-out 2-3, 71-4
Remote Access
log-off user 71-30
settings page 11-36
tables
showing/hiding columns 71-8
sorting columns 71-8
trending 71-2
viewing related events in Event Viewer 69-58
views
closing 71-23
custom 71-24
docking 71-24
floating 71-24
list 71-21
opening 71-23
tiling 71-23
HTML file
export HPM data as 71-31
HTTP
Cisco IOS routers
AAA tab 63-32
Command Authorization Override dialog box 63-34
defining policies 63-29
HTTP Policy page 63-31
overview 63-28
Setup tab 63-31
PIX/ASA/FWSM 49-2
configuration 49-3
HTTP (ASA, PIX) class map objects
creating 17-22
HTTP (ASA7.1.x/PIX7.1.x/FWSM3.x/IOS) policy map objects
creating 17-22
properties 17-56
HTTP (ASA7.2+/PIX7.2+) policy map objects
creating 17-22
properties 17-64
HTTP (IOS) class map objects
creating 21-16
creating for zone-based firewall content filtering 21-36
match criteria 21-22
HTTP (Zone Based IOS) policy map objects
creating 21-16, 21-36
match conditions and actions 21-35
HTTP class map objects
match criteria 17-66
HTTP-FORM
settings in AAA server objects 6-44
HTTP policy
overriding HTTPS port number 3-47
sharing
HTTPS port number 3-47
HTTP policy map objects
match conditions and actions 17-66
HTTP proxy server
configuring for IPS global correlation 36-24
HTTP Response Code 500 deployment errors 9-15
HTTPS
setting up 2-3
troubleshooting certificate errors 9-5
hub-and-spoke topology
description 25-2
joined hub-and-spoke topology 25-5
tiered hub-and-spoke topologies 25-5

ICMP rules
PIX/ASA/FWSM 49-4
add/edit 49-5
ICMP settings
configuring on IOS routers 62-18
icons
Configuration Manager toolbar reference 1-39
event table toolbar reference 69-16
Event Viewer status color code 69-31
map elements 35-14
ICQ class map objects
creating 21-16
match criteria 21-21
identity-aware firewall policies
collecting user statistics 13-25
configuring 13-7
configuring cut-through proxy 13-23
configuring identity options 13-15
configuring layer 2 SGT imposition 46-44
configuring rules 13-21
configuring security group tagging 46-44
configuring the ASA 13-7, 14-8
enabling 13-8
filtering VPN traffic 13-26
identifying AD servers and agents 11-38, 13-8
managing 13-1
monitoring 13-27
overview 13-1
requirements 13-3
user identity acquisition 13-2
Identity Configuration wizard
Active Directory Agent Settings 13-13
Active Directory Settings 13-11
Preview 13-15
Identity Settings page 11-38
identity user group objects
creating 13-19
selecting 13-21
user identity acquisition 13-2
idle timeout, Security Manager client 11-10
IDM
device manager 72-15
IDSM
adding when using non-default HTTPS (SSL) port 3-7
Create and Edit IDSM Data Port VLANs dialog boxes 68-49
Create and Edit IDSM EtherChannel VLANs dialog boxes 68-49
credentials 3-19
defining Data Port VLANs 68-46
defining EtherChannel VLANs 68-44
deleting Data Port VLANs 68-47
deleting EtherChannel VLANs 68-45
deployment failures when changing data port VLAN running mode 9-17
IDSM Settings page 68-47
IDSM Slot-Port Selector dialog box 68-50
mode support limitations 68-43
troubleshooting deployment 9-16
understanding settings on Catalyst devices 68-43
IE 10 security settings 10-2
IGMP
PIX/ASA/FWSM
Access Group parameters 55-5
Access Group tab 55-5
enable 55-1
Join Group parameters 55-7
Join Group tab 55-7
page 55-2
parameters 55-4
Protocol tab 55-3
Static Group parameters 55-6
Static Group tab 55-6
ignore error message, configure Security Manager to 9-10
IKE (Internet Key Exchange)
comparing version 1 and 2 26-4

configuring IKE and IPsec policies 26-1
configuring IKEv2 authentication 26-68
configuring proposal 26-9
Diffie-Hellman modulus groups 26-7
encryption algorithms 26-6
hash algorithms 26-6
IKEv2 Authentication policy 26-70, 26-72
overview 26-2
selecting the IKE version for devices in site to site VPNs 26-26
understanding 26-5
IKE keepalive
understanding 26-33
IKE proposal objects
v1 properties 26-10
v2 properties 26-14
IKE proposals (policies)
in GET VPNs 29-15
IKEv2 Authentication dialog box 26-72
IKEv2 Authentication page 26-70
IKEv2 settings
configuring 26-37
configuring cookie challenges 26-37
IM (ASA7.2+/PIX7.2+) policy map objects
creating 17-22
properties 17-70
IM (IOS) policy map objects
creating 17-22
properties 17-73
IM (Zone Based IOS) policy map objects
creating 21-16
match conditions and actions 21-35
IM (Zone based IOS) policy map objects
creating 21-16
Image Management 73-1
supported versions 73-2
Image Manager 73-9, 73-16
abort installation job 73-35
Add Image 73-11
Bootstrapping Devices 73-8
bundled images 73-30
bundles 73-13
create 73-13
delete 73-15
rename 73-15
view images 73-14
compatible images 73-17
configuring install location 73-19
device memory 73-18
devices 73-16
Getting Started 73-1
Installation Job Summary 73-33
installation wizard 73-26
installing compatible images on devices 73-30
installing images on selected devices 73-31
job approval workflow 73-36
jobs 73-32
RAM 73-17
Repository 73-9
retry on installation failure 73-35
roll back 73-35
settings 11-41
supported image types 73-5
supported platforms 73-2
Troubleshooting 73-37
update validation 73-23
updating images on devices 73-20
Using 73-1
Admin Settings 73-6
View All Images 73-10
view device information 73-16
view installation job details 73-34
Image Manager command 1-38
images
view 73-10
image updates 73-20
IMAP
configuring for inspection rules 17-21

IMAP class map objects
creating 21-16
match criteria 21-25
IM applications
match conditions for zone-based firewalls 21-21
protocol information for IM application inspection 21-33
IMAP policy map objects
creating 21-16
match conditions and actions 21-35
IM class map objects
creating 17-22
match criteria 17-71
IM policy map objects
match conditions and actions 17-71
import
device inventory 3-31
device with policies 10-13
policy objects 6-23
Import Background Image dialog box 35-13
Import Rules wizard
Enter Parameters page 16-41
Preview page 16-43
Status page 16-42
inheritance
inheriting rules 5-47
understanding 5-4
understanding signature policies 39-3
versus assignment 5-6
Inherit Rules command 1-32
Inherit Rules dialog box 5-47
Inspect/Application FW Rule wizard
Address and Port page 17-13
Inspected Protocol page 17-17
Match Traffic page 17-11
inspection
deny rules 17-5
global correlation (IPS)
configuring 42-5
inspection map objects
understanding 6-78
inspection rules
ACL naming conventions 12-5
add/edit rule wizard 17-11, 17-13, 17-17
choosing interfaces 17-2
configuring 17-5
configuring custom protocol name 17-22
configuring DNS settings 17-19
configuring ESMTP settings 17-20
configuring fragment inspection 17-20
configuring identity aware 13-21
configuring in Map view 35-23
configuring RPC settings 17-21
configuring security group aware 14-17
configuring settings for IOS devices 17-111
configuring settings in Map view 35-24
configuring SMTP settings 17-20
deep inspection options
IMAP 17-21
POP3 17-21
deleting 12-9
disabling 12-20
editing 12-10
enabling 12-20
Inspection Rules page 17-8
managing 17-1
moving 12-19
preserving ACL names 12-4
preventing DoS attacks on IOS devices 17-5
selecting protocols 17-3, 17-17
understanding 17-2
understanding access rule requirements 17-4
understanding NAT effects 12-3
understanding processing order 12-2
Inspection Rules page 17-8
Inspection settings page 17-111
inspect maps
policy maps

Add Country Network Codes dialog box 17-48
Edit Country Network Codes dialog box 17-48
Inspect parameter map objects
properties 21-31
Inspect Parameters map objects
creating 21-16, 21-36
installing
Security Manager client 1-12
Integrated Local Management Interface (ILMI) 62-50
Interactive Authentication Configuration dialog box 15-24
Interface Name Conflict dialog box 6-78
Interface Properties dialog box 35-19
Interface Role Contents dialog box 12-14
interface role objects
creating 6-74
defining subinterfaces 6-76
distinguishing from interfaces 6-76
handling conflicts between role and interface names 6-78
Interface Role dialog box 6-75
specifying during policy definition 6-76
understanding 6-73
use when a single interface name is allowed 6-77
interfaces
adding or changing modules 3-40
ASA
edit EtherChannel-assigned interface 46-12
EtherChannels 46-9, 46-13
LACP 46-12
ASA/FWSM
IPv6 46-47, 46-73
IPv6, add/edit 46-52
IPv6, add/edit prefixes 46-54, 46-56
ASA 5505 46-6
ASA devices
Advanced tab 46-41
IP Type 46-58
Catalyst switches and 7600 Series routers
Access Port Selector dialog box 68-30
Create and Edit Interface dialog boxes-Access Port mode 68-9
Create and Edit Interface dialog boxes-Dynamic Port mode 68-18
Create and Edit Interface dialog boxes-Other mode 68-24
Create and Edit Interface dialog boxes-Routed Port mode 68-12
Create and Edit Interface dialog boxes-subinterfaces 68-22
Create and Edit Interface dialog boxes-Trunk Port mode 68-14
Create and Edit VLAN dialog boxes 68-28
Create and Edit VLAN Group dialog boxes 68-34
defining ports 68-6
deleting ports 68-7
generating names 68-6
Interfaces/VLANs page-Interfaces tab 68-8
Interfaces/VLANs page-Summary tab 68-3
Interfaces/VLANs page-VLAN Groups tab 68-33
Interfaces/VLANs page-VLANs tab 68-27
Service Module Slot Selector dialog box 68-35
Trunk Port Selector dialog box 68-31
understanding 68-5
VLAN Selector dialog box 68-35
Cisco IOS routers
Advanced Interface Settings dialog box 62-16
Advanced Interface Settings page 62-16
available types 62-2
Create Router Interface dialog box 62-8
defining advanced settings 62-13
defining basic settings 62-4
defining CEF interface settings 62-25
defining IPS module settings 62-22
deleting from 62-6
generating names 62-4
Interface Auto Name Generator dialog box 62-12
overview 62-1
Router Interfaces page 62-7

understanding helper addresses 62-14
configuring IOS IPS rules 45-9
configuring multiple contexts 59-3
distinguishing from interface roles 6-76
failover
MAC address 50-23
PIX/ASA/FWSM 50-23
PIX 6.3 50-12
IPS
configuring 37-6
configuring bypass mode 37-12
configuring CDP mode 37-12
configuring inline interface pairs 37-13
configuring inline VLAN pairs 37-14
configuring physical 37-9
configuring VLAN groups 37-15
deploying VLAN groups 37-5
inline interface mode 37-3
inline VLAN pair mode 37-3
interfaces policy 37-6
managing interface configurations 37-1
physical interface properties 37-10
promiscuous mode 37-2
roles 37-1
sensing modes overview 37-2
understanding 37-1
viewing summary 37-8
VLAN group mode 37-4
IP Type
PIX 6.3 46-30
PIX/ASA
allocation in security contexts 59-12
IP Type 46-58
PPPoE Users 46-71
redundant 46-8
subinterfaces 46-7, 46-15
VPDN groups 46-72
PIX/ASA/FWSM
add/edit 46-31
Advanced settings 46-68
configuring 46-3
contexts 46-5
DDNS update rules 52-19
enabling traffic between same security levels 46-70
General tab 46-33
manage 46-26
management access 49-6
understanding 46-3
PIX/ASA 7+ devices
MAC address 46-60
PIX 6.3
add/edit 46-28
routed and transparent 46-5
specifying during policy definition 6-76
specifying subinterfaces 6-76
throughput delay 62-18
Interface Selector dialog box (VLAN ACL Content) 68-42
Interfaces page (IPS) 37-6
Interface Specific Authentication Server Groups dialog box 31-16
Interface Specific Client Address Pools dialog box 31-12
inventory
deleting devices from 3-58
export devices
DCR, CS-MARS, Security Manager formats 10-6
device with policies 10-6
overview 10-6
supported CSV formats 10-9
using command line utility 10-10
import devices
device with policies 10-13
inventory, device
adding devices 3-6
adding devices from configuration files 3-22
adding devices from inventory file 3-31
adding devices from network 3-12
adding devices manually 3-26

device status view
working with 3-64
managing 3-1
testing device connectivity 9-1
troubleshooting device discovery failures 3-7
understanding 3-1
understanding contents 3-3
understanding device clusters 3-9
understanding generic devices 3-8
viewing inventory status 72-12
working with 3-36
Inventory Status command 1-35
Inventory Status window 72-13
Inverse ARP 62-61
inverse multiplexing over ATM (IMA) 62-40
IOS devices
configuring transparent firewall rules 23-1
remote access IPSec VPNs
user group policies 33-13
remote access IPsec VPNs
creating using wizard 30-36
remote access SSL VPNs
configuring bookmarks 31-82
configuring WINS servers for file system access 31-88
creating using wizard 30-32
remote access VPNs
configuring SSL VPN policies 33-14
Context Editor dialog box (IOS) 33-15, 33-16
Dynamic VTI/VRF Aware IPsec settings 33-7
high availability 33-11
IPsec proposals 33-4
SDM 72-16
IOS IPS
affect of load balancing 45-8
comparing to IPS appliances and service modules 36-2
configuration files 45-3
configuration overview 45-4
configuring 45-1
configuring general settings 45-7
configuring interface rules 45-9
configuring target value ratings 40-17
event actions
filter rule attributes 40-9
filter rules 40-4, 40-7
filter rules tips 40-6
network information 40-17
overrides 40-13
overview 40-1
possible actions 40-2
process overview 40-1
settings 40-23
getting started 36-1
initial preparation of router 45-5
lightweight signature engines 45-2
limitations and restrictions 45-3
selecting signature category 45-6
signatures
adding custom 39-19
cloning 39-21
configuring 39-4
defining 39-1
detailed information 39-2
editing 39-14
editing Meta engine component list 39-29
editing or tuning parameters 39-23
enabling or disabling 39-14
engines 39-20
exporting 39-9
inheritance 39-3
parameters list 39-24
policy 39-4
shortcut menu 39-10
understanding 39-1
viewing update level 39-9, 39-13
understanding 45-1
understanding subsystems and revisions 45-2

IOS Software Release 12.1 and 12.2
managing routers 61-3
IOS Web Filter Exclusive Domain Name dialog box 18-14
IOS Web Filter Rule and Applet Scanner dialog box 18-13
IP address
supporting dynamic 3-36
IP addresses
network masks 6-81
specifying in policies 6-87
IP Intelligence
settings 11-41
IP Intelligence dialog box 72-35
IP Intelligence in Report Manager 72-35
IP Intelligence Settings in Dashboard 72-2
IP Intelligence using Quick Launch 72-35
IP Intelligence widget 72-35
IP Options policy map objects
creating 17-22
properties 17-75
IPS
IPS Module router interface settings policies 62-22
MPC rule wizard
tab 58-8
PIX/ASA/FWSM
identity-aware rules 13-21
rules 58-5
IPS alerts
properties 69-18
IPS Certificates dialog box 44-10
IPS command 1-34
IPS Devices
selecting for Event Viewer 69-34
IPS devices
adding SSL thumbprints manually 9-5
allowed hosts 36-7
anomaly detection
configuring 41-6
configuring histograms 41-11
configuring learning accept mode 41-8
configuring signatures 41-4
configuring thresholds 41-11
detection zones 41-3
managing 41-1
modes 41-2
understanding 41-1
understanding histograms 41-9
understanding thresholds 41-9
understanding worms 41-2
when to turn off 41-4
blocking
configuring 43-7
configuring ARC 43-1
configuring blocking devices 43-14
configuring master blocking sensors 43-13
configuring never block hosts and networks 43-17
configuring router blocking interfaces 43-15
configuring user profiles 43-12
configuring VLAN blocking interfaces 43-16
general options 43-10
master blocking sensor 43-6
policy 43-8
rate limiting 43-4
router and switch blocking devices 43-4
strategies 43-3
understanding 43-1
capturing network traffic 36-2
certificates 44-10
changing those selected for reports 70-22
configuration overview 36-5
configuration overview for IOS IPS 45-4
configuring AAA 36-21
configuring Analysis Engine global variables 36-30
configuring DNS servers 36-24
configuring for event management 69-29
configuring for report management 70-3
configuring HTTP proxy server 36-24
configuring NTP 36-23

configuring OS maps 40-21
configuring SNMP 36-8
configuring target value ratings 40-17
configuring the external product interface 36-26
configuring user accounts 36-18
credentials, IPS router modules 3-20
deployment of passwords 36-17
deployment topology 36-4
discovery of passwords 36-17
event actions
  example filter rule 69-67
  filter rule attributes 40-9
  filter rules 40-4, 40-7
  filter rules tips 40-6
  network information 40-17
  overrides 40-13
  overview 40-1
  possible actions 40-2
  process overview 40-1
  settings 40-23
Event Viewer support 69-4
getting started 36-1
global correlation
  configuring 42-1
  configuring inspection and reputation 42-5
  configuring network participation 42-7
  data collected 42-3
  requirements and limitations 42-4
  understanding 42-1
  understanding network participation 42-3
  understanding reputation 42-2
initializing 2-10
interfaces
  configuring 37-6
  configuring bypass mode 37-12
  configuring CDP mode 37-12
  configuring inline interface pairs 37-13
  configuring inline VLAN pairs 37-14
  configuring physical 37-9
  configuring VLAN groups 37-15
  deploying VLAN groups 37-5
  inline interface mode 37-3
  inline VLAN pair mode 37-3
  interfaces policy 37-6
  managing interface configurations 37-1
  physical interface properties 37-10
  promiscuous mode 37-2
  roles 37-1
  sensing modes overview 37-2
  understanding 37-1
  viewing summary 37-8
  VLAN group mode 37-4
IPS modules for ASA 58-15
license, exporting 11-59
licenses
  automating 44-3
  managing 44-1
  redeploying 44-2
  updating 44-1
looking up signature policies for CS-MARS events 72-45
looking up signature policies for Event Viewer events 69-54
managing 44-1
managing user accounts and passwords 36-15
monitoring
  removing false positive IPS events 69-66
passive OS fingerprinting 40-19
password requirements 36-20
policy discovery 5-14
rebooting 44-12
Report Manager reports
  general VPN reports 70-19
  IPS top reports 70-17
rollback restrictions 8-66
showing containment 3-56
signatures
  adding custom 39-19

cloning 39-21
configuring 39-4
configuring settings 39-30
defining 39-1
detailed information 39-2
editing 39-14
editing Meta engine component list 39-29
editing or tuning parameters 39-23
enabling or disabling 39-14
engines 39-20
exporting 39-9
inheritance 39-3
parameters list 39-24
policy 39-4
shortcut menu 39-10
understanding 39-1
viewing update level 39-9, 39-13
SSL certificate configuration 11-22
traffic flow notifications 36-30
tuning recommendations 36-4
understanding managed and unmanaged passwords 36-16
understanding network sensing 36-2
understanding user roles 36-15
updates
automatically applying 44-6
checking for and downloading 44-5
configuring server 44-4
managing 44-4
manually applying 44-7
user account attributes 36-19
viewing signature events in CS-MARS 72-44
viewing signature events in Event Viewer 69-57
virtual sensors
advantages 38-3
assigning interfaces 38-4
attributes 38-7
configuring 38-1, 38-5
deleting 38-10
editing policies 38-9
identifying 38-5
inline TCP session tracking mode 38-3
Normalizer mode 38-4
renaming 38-8
restrictions 38-3
understanding 38-1
IPsec
remote access VPNs
access policies for IKEv2 (ASA), configuring 31-49
access policies for IKEv2 (ASA), reference 31-45
access policies for IKEv2 (ASA), understanding 31-44
certificate to connection profile map policy (IKEv1) 31-36
certificate to connection profile map rules (IKEv1) 31-37
cluster load balancing 31-5
configuring IKE and IPsec policies 26-1
connection profiles 31-7
connection profiles (ASA, PIX 7+) 31-8
creating on ASA/PIX 7.0+ 30-25
creating on IOS/PIX 6.3+ 30-36
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
Dynamic VTI/VRF Aware IPsec settings 33-7
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
high availability policies 33-11
IKE proposals 26-9
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
NAT settings 26-42
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9

understanding 30-2
understanding IKE 26-5
understanding NAT settings 26-41
user group policies 33-13
VPNSM, VPN SPA, VSPA settings 33-6
wizard 30-13
IPsec/GRE VPN
advantages of IPsec tunneling with GRE 27-3
configuring 27-5
configuring GRE modes 27-6
dynamically addressed spokes 27-5
implementation 27-3
overview 27-1, 27-2
prerequisites for successful configuration 27-3
supported platforms 25-9
understanding 27-2
IPSec Client Software Update dialog box 31-22
IPsec Pass Through policy map objects
creating 17-22
properties 17-80
IPsec Proposal Editor dialog box
ASA and PIX 7.0+ devices 31-41
IOS and PIX 6.3 devices 33-4
IPsec proposals
configuring for Easy VPN 28-10
configuring for remote access VPNs
attributes for ASA and PIX 7.0+ devices 31-41
attributes for IOS and PIX 6.3 devices 33-4
configuring in site-to-site VPNs 26-22
overview 26-2
remote access VPNs
attributes for ASA and PIX 7.0+ devices 31-41
attributes for IOS and PIX 6.3 devices 33-4
configuring for ASA and PIX 7.0+ devices 31-40
configuring for IOS and PIX 6.3 devices 33-3
selecting the IKE version for devices 26-26
understanding 26-18
understanding crypto maps 26-19
understanding site-to-site 26-19
understanding transform sets 26-20
using reverse route injection 26-21
IPsec technologies
defining 25-30
mandatory and optional policies 25-6
policies 25-5
supported platforms 25-9
supported platforms for remote access VPNs 30-8
understanding 25-5
IPSec transform set objects
attributes 26-27
understanding 26-20
IPSec VPN
zone-based firewalls 21-6
IPS event
definition of 40-1
IPS Health Monitor page in Dashboard 72-2
IPS interfaces
IPS Monitoring Information dialog box 62-24
IPS module
credentials 3-20
IPS Module Discovery dialog box 3-20
IPS Module interface settings policies 62-22
IPS Rules dialog box 45-10
IPS sensor
IDM 72-15
IPS sensors
default transport protocol 11-22
IPS signatures
finding from CS-MARS events 72-45
finding from Event Viewer events 69-54
tuning 69-66
viewing related CS-MARS events 72-44
viewing related events in Event Viewer 69-57
IPS tab, Licensing page 11-58
IPS Updates page 11-47
IP Type
interface configuration
ASA and PIX 7+ 46-58

PIX 6.3 46-30
IPv4 pool objects
attributes 6-92
IPv6
interfaces
add/edit 46-52
add/edit prefixes 46-54, 46-56
ASA/FWSM 46-47, 46-73
management IPv4 address requirements 1-8
Neighbor cache 47-7
specifying addresses in policies 6-87
support in Security Manager 1-8
IPv6 access rules
ACL naming conventions 12-5
deleting 12-9
disabling 12-20
editing 12-10
enabling 12-20
expiration dates 16-22
identity-aware rules
requirements 13-3
moving 12-19
preserving ACL names 12-4
sharing ACLs among interfaces 11-18
understanding global 16-3
understanding processing order 12-2
IPv6 policy map objects
match conditions and actions 17-78, 17-92
properties 17-77, 17-91
IPv6 pool objects
attributes 6-93
IPv6 static routes
PIX/ASA/FWSM
configuration 56-131
ISAKMP/IPsec settings
configuring 26-33
ISE Settings page 11-56
ISR
zone-based firewall
restrictions 21-3

J

job deployment methods
understanding 8-8
jobs
aborting 8-55
approving 8-39
creating and editing deployment in non-Workflow mode 8-28
creating and editing deployment in Workflow mode 8-35
Deployment Manager 8-16
discarding 8-41
including devices in 8-8
rejecting 8-39
states
Workflow mode 8-6
submitting 8-38
joined hub-and-spoke topology 25-5
Join Group tab (IGMP) 55-7
JumpStart 1-24
Jumpstart command 1-38

K

Kazaa2 class map objects
creating 21-16
match criteria 21-21
Kerberos
configuring constrained delegation (KCD) 31-69
description 6-29
settings in AAA server objects 6-39
understanding constrained delegation (KCD) 31-66
key encryption key (KEK), GET VPN 29-4
key servers
adding 29-19
choosing the rekey transport mechanism 29-6

communication flow 29-2
cooperative, for redundancy 29-7
editing 29-19
generating, synchronizing RSA keys 29-13
registration failures 29-8
registration process 29-4
security policy ACLs 29-10
key servers (GET VPN)
configuring 29-18
Key Servers page (GET VPN) 29-18
Key Servers Selection dialog box 29-21
knowledge base structure (IPS) 41-8

L

LACP
interface assigned to an EtherChannel 46-12
large scale Dynamic Multipoint VPN (DMVPN)
mandatory and optional policies 25-6
Launch menu 1-37
Report Manager 70-8
layer 2 SGT imposition 46-44
LDAP
settings in AAA server objects 6-40
LDAP Attribute Map objects
attributes 6-46
learning accept mode (IPS), configuring 41-8
licenses
configuring for ASA devices 2-9
configuring for IOS devices 2-10
exporting IPS 11-59
IPS
automating 44-3
managing 44-1
redeploying 44-2
updating 44-1
Security Manager 10-16
License Update Status Details dialog box 11-62
licensing
Settings page 11-57
Lightweight Directory Access Protocol (LDAP)
description 6-29
lightweight signature engines 45-2
line access
Cisco IOS routers
Console Policy page 63-42
overview 63-35
VTY Policy page 63-50
Link Aggregation Control Protocol 46-12
Link Properties dialog box 35-20
load balancing
configuring in large scale DMVPN 27-16, 27-17
configuring IOS IPS deny actions 45-8
server attributes in large scale DMVPN 27-17
Local Policy Will Be Replaced dialog box 5-44
Local Web Filter class map objects
match criteria 21-29
Local web filter class map objects
creating 21-36
Local Web Filter parameter map objects
properties 21-38
Local web filter parameter map objects
creating 21-36
locking
activities 4-3
devices and policies 5-9
objects 5-10
understanding 5-8
VPN topologies 5-10
Log Buffer window 72-18
logging
Cisco IOS routers
defining NetFlow interfaces 65-15
defining NetFlow parameters 65-6
defining syslog servers 65-3
Logging Setup Policy page 65-7
NetFlow policy page 65-12
overview 65-1

Syslog Server dialog box 65-11
Syslog Servers Policy page 65-10
syslog setup parameters 65-1
syslog severity levels 65-4
PIX/ASA/FWSM 54-1
email notifications 54-8
email recipients 54-8
embedded event manager 54-3
event lists 54-9
event lists, add/edit 54-10
filters 54-12
filters, editing 54-13
levels 54-24
logging setup 54-14
message classes and IDs 54-9
message editing 54-25
message limits 54-18
message limits, add/edit 54-18
NetFlow 54-1
NetFlow, add/edit collector 54-2
rate limit levels 54-17
rate limits, add/edit 54-19
server 54-21
server setup 54-20
set-up 54-15
syslog class 54-11
syslog message ID 54-11
syslog servers 54-26, 54-27
syslog servers, add/edit 54-28
syslog messages supported for CS-MARS queries 72-46
logging in to
Cisco Security Management Suite server 1-12
CiscoWorks Common Services 1-12
logging into
Security Manager 1-11, 1-12
Logging page, IPS platform 36-30
logs
configuring audit log default settings 11-62
configuring debug levels 11-11
Logs page 11-62
loopback cells 62-51
low-latency queuing (LLQ) 66-5

M

MAC address
interface configuration
ASA and PIX 7+ 46-60
PIX/ASA/FWSM
add/edit 47-8
interface 50-23
learning 47-9
learning, enable/disable 47-9
table 47-8
MAC address pool objects
attributes 6-94
MAC exempt lists
configuring 15-7, 15-26
rule attributes 15-27
Maintenance Operation Protocol (MOP), enabling 62-19
Management Access
PIX/ASA/FWSM
interface 49-6
management address
requirements for IPv6 devices 1-8
Management Center for Cisco Security Agents
configuring connection to IPS devices 36-26
connection attributes 36-27
posture ACLs 36-29
Management IP address
PIX/ASA/FWSM 47-10
Management IPv6
ASA 5505 47-11
Manage menu 1-34
Map menu 1-33
map objects
class maps

creating for inspection rules 17-22
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
parameter maps
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
Inspect properties 21-31
Local Web Filter properties 21-38
N2H2 properties 21-39
Protocol Info properties 21-33
Trend properties 21-42
URLF Glob properties 21-45
URL Filter properties 21-43
Websense properties 21-39
policy maps
creating for inspection rules 17-22
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
DCE/RPC properties 17-29
DNS properties 17-32
ESMTP properties 17-39
FTP properties 17-42
GTP properties 17-45
H.323 (ASA/PIX/FWSM) properties 17-51
HTTP (ASA7.1.x/PIX7.1.x/FWSM3.x/IOS) properties 17-56
HTTP (ASA7.2+/PIX7.2+) properties 17-64
IM (ASA7.2+/PIX7.2+) properties 17-70
IM (IOS) properties 17-73
IP Options properties 17-75
IPsec Pass Through properties 17-80
IPv6 properties 17-77, 17-91
NetBIOS properties 17-81
regular expression group properties 17-108
regular expression properties 17-108
SIP (ASA/PIX/FWSM) properties 17-83, 17-93, 17-102, 17-103
Skinny properties 17-87
SNMP properties 17-90
TCP Map properties 58-22
Web Filter properties 21-47
regular expression objects
metacharacters 17-109
understanding 6-78
Map Properties command 1-33
Map Rule dialog box
connection profile map matching rules 31-39
connection profile maps 31-39
maps
access permissions 35-8
adding existing managed devices 35-16
adding new managed devices 35-16
arranging elements 35-11
background color 35-13
background images
deleting 35-13
importing 35-13
scale and position 35-13
setting 35-13
centering elements 35-11
changing the zoom level 35-11
class maps
Class Map dialog box 17-28, 21-19
creating 35-9
default map 35-9
deleting 35-10
displaying devices from Device View 35-16
displaying managed devices 35-16
displaying your network 35-14
elements, understanding 35-14
excluding private and reserved networks 11-3
exporting 35-11
icons 35-14
layer 3 links
autolink settings 11-3
creating 35-19

deleting 35-19
layouts, using 35-11
linking maps 35-13
navigation window 35-4
objects
adding 35-17
deleting 35-17
opening 35-10
overview 35-1
panning 35-11
refreshing 35-1
removing managed devices 35-16
renaming 35-10
saving 35-10
searching for nodes 35-12
selecting elements 35-12
setting background 35-13
showing containment for Catalyst, ASA, PIX, IPS devices 35-16
understanding 35-1
undocking window 35-2
working with 35-8
Map Settings dialog box 35-13
Map View
cloning devices 35-22
configuring firewall policies 35-23
configuring firewall settings policies 35-23
context menu
Layer 3 link 35-7
managed device node 35-5
map background 35-7
map objects 35-7
selected nodes 35-6
VPN connection 35-6
device policies, managing 35-22
discovering device configurations 35-22
icons for elements 35-14
main page 35-2
menus, context 35-5
navigation window 35-4
performing basic policy management 35-22
previewing device configurations 35-22
sharing device policies 35-22
toolbar reference 35-4
VPNs
creating 35-21
displaying existing 35-21
editing or showing peers 35-22
editing policies 35-22
managing 35-20
Map view
Autolink Settings page 11-3
copying between devices 35-22
overview 1-18, 35-1
Map View command 1-32
master blocking sensor 43-6
Master Blocking Sensor dialog box 43-13
maximum receive reconstructed unit (MRRU) 62-82
maximum segment size (MSS) 62-17
MBoundary
PIX/ASA/FWSM
configuration 55-9
interface configuration 55-10
MD5 hash algorithm 26-7
memory-allocation lite 63-80
memory settings
Cisco IOS routers
defining 63-78
overview 63-78
Memory Policy page 63-79
menu reference
Activities 1-36
Configuration Manager overview 1-29
Edit (Configuration Manager) 1-31
File (Configuration Manager) 1-30
File (Event Viewer) 69-9
File (Report Manager) 70-8
Help (Configuration Manager) 1-38

Launch 1-37
Launch (Report Manager) 70-8
Manage 1-34
Map 1-33
Policy (Configuration Manager) 1-32
Tickets 1-36
Tools (Configuration Manager) 1-34
Tools (Report Manager) 70-8
View (Configuration Manager) 1-31
View (Event Viewer) 69-10
message
editing
PIX/ASA/FWSM 54-25
PIX/ASA/FWSM
limits 54-18
limits, add/edit 54-18
rate limits, add/edit 54-19
message classes and IDs
PIX/ASA/FWSM 54-9
metacharacters
URLF Glob parameter maps 21-46
Mobile application for CSM 72-11
Modify Access List dialog box (Allowed Hosts policy) 36-7
Modify Physical Interface Map dialog box 37-10
monitoring
CS-MARS
integrating with Security Manager 72-36
device managers, using 72-14
device status 72-1
network activities 72-1
PRSM, launching 72-20
monitoring widget for server 72-7
mount point
PIX/ASA
add/edit 48-19, 48-20
mount point configuration
ASA 48-18
Move Row Down command 1-31
Move Row Up command 1-31
MPC
a.k.a. Modular Policy Framework 58-6
MRoute
PIX/ASA/FWSM
configuration 55-8
MRoute page
description 55-8
MSN Messenger class map objects
creating 21-16
match criteria 21-21
multicast
PIX/ASA/FWSM
Enable PIM and IGMP 55-1
IGMP Access Group parameters 55-5
IGMP Access Group tab 55-5
IGMP Join Group parameters 55-7
IGMP Join Group tab 55-7
IGMP parameters 55-4
IGMP Protocol tab 55-3
IGMP Static Group parameters 55-6
IGMP Static Group tab 55-6
MBoundary configuration 55-9
MBoundary interface configuration 55-10
MRoute configuration 55-8
Multicast Boundary Filter page 55-9
Multicast Group, add/edit 55-19, 55-21
Multicast Group rule 55-17
PIM Bidirectional Neighbor Filter 55-14
PIM Bidirectional Neighbor Filter tab 55-13
PIM Neighbor Filter 55-13
PIM Neighbor Filter tab 55-12
PIM page 55-11
PIM Protocol dialog box 55-12
PIM Protocol tab 55-11
PIM Rendezvous Point, add/edit 55-16
PIM Rendezvous Points tab 55-15
PIM Request Filter tab 55-18, 55-20
PIM Route Tree tab 55-17

Multicast Boundary Filter page
description 55-9
multicast rekey in GET VPN 29-6
multicast routing
PIX/ASA/FWSM
configuring on 55-1
IGMP 55-2
multicast boundary filters 55-9
multicast routes 55-8
PIM 55-11
Multiclass Multilink PPP (MCMP) 62-75
multilink PPP (MLP) 62-71
defining bundles 62-75
multiple users
activities 4-4
tickets 4-4

N

N2H2 (Smartfilter)
configuring for web filter rules policies 18-15, 18-19
configuring for zone based firewall rules policies 21-36, 21-39, 21-41
N2H2 class map objects
creating 21-36
match criteria 21-30
N2H2 parameter map objects
creating 21-36
properties 21-39
NAC
posture validation not occurring 9-15
NAT
VPN traffic sent unencrypted 9-15
NAT policies
Add/Edit Per-Session NAT rules dialog boxes 24-47
NBAR
enabling protocol discovery 62-19
Neighbor cache 47-7
Neighbor Filter
PIM
PIX/ASA/FWSM 55-13
Neighbor Filter tab
PIM 55-12
NetBIOS logout probe
configuring 13-15
requirements 13-5
NetBIOS policy map objects
creating 17-22
properties 17-81
NetFlow
Cisco IOS routers 65-1, 65-5
interface settings 65-15
configuring
on Cisco IOS routers 65-6
CS-MARS query 72-47
IOS routers 65-12
PIX/ASA/FWSM 54-1
add/edit collector 54-2
network/host objects
attributes 6-83
attributes, NAT 24-42
creating 6-82
naming when provisioned as object groups 6-107
network masks 6-81
optimizing when deploying firewall rules 12-35
understanding 6-80
unspecified value objects 6-86
using in Event Viewer filters 69-67
network access device (NAD) 64-9
Network Address Translation (NAT)
Add/Edit Per-Session NAT rules dialog boxes 24-47
ASA 8.3+
Add/Edit NAT rules dialog boxes 24-36
Translation Rules page 24-34
understanding 24-4
ASA 8.3 devices 24-33
Cisco IOS routers 24-5
Dynamic Rule dialog box 24-11

dynamic rules 24-10
Interface Specification 24-6
Static Rule dialog box 24-7
static rules 24-6
Static Rules tab 24-6
timeouts 24-13
configuring global options for VPNs 26-42
non-ASA 8.3 devices 24-18
No Proxy ARP 24-39, 24-45
PAT pool 24-41
Per-session NAT rules 24-46
PIX/ASA/FWSM
Address Pool dialog box 24-19
Address Pools page 24-18
Advanced NAT Options dialog box 24-29
clearing XLATE on deployment 60-1
configuring on 24-15
configuring translation rules 24-19
Dynamic Rules dialog box 24-23
Dynamic Rules tab 24-22
General tab 24-31
non ASA 8.3 24-18
Policy Dynamic Rules dialog box 24-25
Policy Dynamic Rules tab 24-24
Select Address Pool 24-24
Static Rules dialog box 24-27
Static Rules tab 24-26
Translation Exemptions (NAT 0 ACL) dialog box 24-21
Translation Exemptions (NAT 0 ACL) tab 24-20
Translation Options page 24-16, 24-17
Translation Rules page 24-19
translation types 24-3
transparent mode 24-16
understanding 24-2
round robin allocation 24-41
understanding NAT effects on firewall rules 12-3
understanding NAT settings for VPNs 26-41
understanding NAT traversal 26-41
Network Admission Control (NAC)
Cisco Trust Agent 64-9
components 64-9
defining identity parameters 64-13
defining interface parameters 64-11
defining setup parameters 64-10
Identities tab 64-18
Identity Action dialog box 64-19
Identity Profile dialog box 64-19
Interface Configuration dialog box 64-17
Interfaces tab 64-16
NAC Policy page 64-14
network access device (NAD) 64-9
on Cisco IOS routers 64-8
Setup tab 64-14
supported platforms 64-8
understanding system flow 64-9
Network Information page (IPS) 40-17
network masks
discontiguous 6-81
discovering 6-82
displaying 6-82
understanding 6-81
network participation, IPS
configuring 42-7
data collected 42-3
requirements and limitations 42-4
understanding 42-3
understanding global correlation 42-1
understanding reputation 42-2
network sensing
capturing network traffic 36-2
deployment topology 36-4
overview 36-2
tuning recommendations 36-4
Network Time Protocol (NTP)
Cisco IOS routers
creating NTP servers 63-97
NTP Policy page 63-98

NTP Server dialog box 63-99
overview 63-96
Never Block Host dialog box 43-17
Never Block Network dialog box 43-17
New Activity command 1-36
New Device command 1-30
New Device Groups command 1-31
New Device wizard
Choose Method page 3-6
Device Grouping page 3-49
Device Information page - Add Device from File 3-33
Device Information page - Configuration File 3-23
Device Information page - Network 3-14
Device Information page - New Device 3-27
New Map command 1-33
New or Edit CS-MARS Device dialog box 11-8
New Ticket command 1-37
NHRP
DMVPN spoke-to-spoke connections 27-11
Node Properties dialog box 35-18
Non-Workflow mode
viewing
device details 8-26
non-Workflow mode
changing modes 1-28
comparing with Workflow mode 1-22
configuration files
deploying 8-28
previewing 8-44
configurations
rolling back 8-69
creating tickets 4-14
deployment 8-3
deployment jobs
aborting 8-55
Deployment Status Details dialog box 8-32
opening tickets 4-15
taking over another user session 10-23
understanding 1-22
No Proxy ARP
NAT rule 24-39, 24-45
PIX/ASA/FWSM Platform 56-1
notifications, e-mail
configuring SMTP server 1-27
NS Lookup 72-26, 72-29
NT
settings in AAA server objects 6-43
NTP
PIX/ASA/FWSM 52-21
server configuration 52-21
NTP policy, IPS platform 36-23
NTP server
configuring for IPS devices 36-23
null0 56-128

O

object groups
policy discovery 5-14
object group search
ASA 8.3+ devices 16-25
PIX 6.3 devices 16-27
objects
AAA server
HTTP-FORM settings 6-44
Kerberos settings 6-39
LDAP settings 6-40
NT settings 6-43
RADIUS settings 6-35
SDI settings 6-43
TACACS+ settings 6-38
AAA server groups
attributes 6-49
creating 6-48
default server groups on IOS devices 6-31
predefined authentication groups 6-30
understanding 6-27
AAA servers

creating 6-32
supported additional types for ASA/PIX/FWSM 6-28
supported types 6-28
understanding 6-27
access control lists
creating 6-53
extended objects 6-54
standard objects 6-56
unified objects 6-58
web objects 6-57
ASA group policies
client configuration settings 34-6
client firewall attributes 34-7
connection settings 34-33
DNS/WINS settings 34-29, 34-30
hardware client attributes 34-9
IPSec settings 34-10
split tunneling settings 34-31
SSL VPN clientless settings 34-12
SSL VPN full client settings 34-19
SSL VPN settings 34-25
technology settings 34-1
AS paths
properties 56-151
basic procedures 6-9
categories, using 6-13
changes in Security Manager 4.4 1-10
Cisco Secure Desktop configuration
creating 33-18
class map
creating for inspection rules 17-22
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
cloning (duplicating) 6-14
community lists
properties 56-153
configuring for ASA routing policies 56-132
configuring for remote access VPN 34-1
creating 6-9
credentials
attributes 28-9
DCE/RPC policy map
properties 17-29
deleting 6-16
DNS policy map
properties 17-32
editing 6-12
ESMTP policy map
properties 17-39
exporting 6-23
file objects
attributes 34-37
selecting 34-39
FlexConfig
creating text objects 7-32
properties 7-30
property selector 7-34
undefined variables 7-33
FlexConfigs
adding to policies 7-35
changing order in policies 7-35
changing variable values 7-35
configuring 7-25
configuring AAA for administrative introducers 63-84
creating 7-28
previewing CLI 7-35
removing from policies 7-35
system variables 7-7
understanding 7-2
variables 7-5, 7-6
FTP policy map
properties 17-42
generating usage reports 6-15
GTP policy map
properties 17-45

H.323 (ASA/PIX/FWSM) policy map
properties 17-51
HTTP (ASA7.1.x/PIX7.1.x/FWSM3.x/IOS) policy map
properties 17-56
HTTP (ASA7.2+/PIX7.2+) policy map
properties 17-64
identity user group
creating 13-19
selecting 13-21
user identity acquisition 13-2
IKE proposals
v1 properties 26-10
v2 properties 26-14
IM (ASA7.2+/PIX7.2+) policy map
properties 17-70
IM (IOS) policy map
properties 17-73
importing 6-23
Inspect parameter map
properties 21-31
interface roles
creating 6-74
IP Options policy map
properties 17-75
IPsec Pass Through policy map
properties 17-80
IPSec transform sets
attributes 26-27
understanding 26-20
IPv6 policy map
properties 17-77, 17-91
LDAP attribute map objects
attributes 6-46
Local Web Filter parameter map
properties 21-38
locking
effects on activities 4-3
managing 6-1
maps
understanding 6-78
N2H2 parameter map
properties 21-39
NetBIOS policy map
properties 17-81
network/host
optimizing when deploying firewall rules 12-35
understanding 6-80
using in Event Viewer filters 69-67
network/host objects
naming when provisioned as object groups 6-107
networks/hosts
creating 6-82
unspecified value objects 6-86
object selectors 6-2
overrides
allowing 6-18
creating for multiple devices 6-19
creating for single device 6-19
deleting 6-21
managing 6-17
understanding 6-18
overview 1-20
parameter map
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
PKI enrollments
defining CA server properties 26-60
defining certificate attributes 26-66
defining enrollment parameters 26-63
defining trusted CA hierarchy 26-67
properties 26-58
policy lists
properties 56-143
policy map
creating for inspection rules 17-22

creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
port forwarding lists
properties 34-40
port list objects
naming when provisioned as object groups 6-107
port lists
creating 6-100
properties 6-102
prefix lists
properties 56-146, 56-148
Protocol Info parameter map
properties 21-33
provisioning as object groups 6-106
regular expression group policy map
properties 17-108
regular expression objects
metacharacters 17-109
regular expression policy map
properties 17-108
route maps 56-136
creating 56-132
understanding 56-132
security group
creating 14-14
selecting for policies 6-2
service objects
naming when provisioned as object groups 6-107
provisioning as object groups 6-108
services
creating 6-100
single sign-on server
properties 34-42
SIP (ASA/PIX/FWSM) policy map
properties 17-83, 17-93, 17-102, 17-103
Skinny policy map
properties 17-87
SLA monitors
attributes 51-10
configuring 51-9
understanding 51-8
SNMP policy map
properties 17-90
SSL VPN Bookmark
configuring 31-82
post URL method and macro substitutions 31-84
SSL VPN Customization
configuring 31-78
creating custom Logon page 31-82
localizing 31-80
SSL VPN gateway
properties 34-64
SSL VPN smart tunnel auto sign-on list
attributes 34-71
SSL VPN smart tunnel list
attributes 34-66, 34-69
configuring 31-85
TCP Map policy map
properties 58-22
text
creating 7-32
time ranges
attributes for recurring ranges 6-72
configuring 6-71
traffic flow
default inspection traffic 58-20
properties 58-18
Trend parameter map
properties 21-42
TrustSec security group
selecting 14-16
URLF Glob parameter map
properties 21-45
URLF Glob parameter maps
metacharacters 21-46
URL Filter parameter map
properties 21-43

user groups
advanced PIX 6.3 settings 34-82
browser proxy settings 34-87
clientless settings 34-83
client VPN software update (IOS) settings 34-81
DNS/WINS settings 34-77
general settings 34-75
IOS client settings 34-78
IOS Xauth settings 34-80
split tunneling settings (Easy VPN/remote access IPSec VPN) 34-77
SSL VPN connection settings 34-88
SSL VPN full tunnel settings 34-84
SSL VPN split tunneling settings 34-86
technology settings 34-73
thin client settings 34-84
using global search to find specific objects 1-42
viewing details 6-14
Web Filter policy map
properties 21-47
Websense parameter map
properties 21-39
WINS server lists
attributes 34-90
creating 31-88
object selectors 6-2
Object Usage dialog box 6-15
Obsoletes dialog box 39-30
OOB (Out of Band) Changes dialog box 8-48
OOB (out of band changes)
avoiding 8-47
detecting and analyzing 8-45
understanding 8-12
Openable Activities dialog box 4-15
Openable Tickets dialog box 4-15
Open Activity command 1-36
Open command (Report Manager) 70-8
Open Map command 1-33
Open Map dialog box 35-10
Open Ticket command 1-37
OS Identifications tab, IPS Network Information policy 40-21
OS Map dialog box 40-22
OSPF
interaction with NAT 56-75
LSAs 56-75
OSPF interfaces
blocking LSA flooding 67-28
defining on Cisco IOS routers 67-25
disabling MTU mismatch detection 67-27
Interface dialog box 67-31
OSPF Interface Policy page 67-30
understanding
authentication 67-29
cost 67-26
network types 67-29
priority 67-26
timer settings 67-28
OSPF parameters
dead interval 56-100, 56-116
hello interval 56-99
hello multiplier 56-100
retransmit interval 56-100, 56-116
transmit delay 56-100, 56-116
OSPF redistribution
defining mappings 67-22
defining maximum prefix values 67-24
understanding 67-22
OSPF routing
Cisco IOS routers
Area dialog box 67-37
Area tab 67-36
defining area settings 67-21
defining interface settings 67-25
defining setup parameters 67-20
Edit Interfaces dialog box 67-36
Max Prefix Mapping dialog box 67-41
OSPF Process Policy page 67-34

overview 67-19
redistributing routes 67-22
Redistribution Mapping dialog box 67-39
Redistribution tab 67-38
Setup dialog box 67-35
Setup tab 67-35
PIX/ASA/FWSM
advanced settings 56-77
Area/Area networks 56-82
Area Range 56-84
Area tab 56-81
Filtering configuration 56-93
Filtering tab 56-92
Filter Rule configuration 56-94
Filter Rule tab 56-94
General tab 56-76
Interface configuration 56-98
Interface tab 56-96
Neighbors tab 56-85
policy 56-75
Range tab 56-84
Redistribution rule 56-87
Redistribution tab 56-86
static neighbor 56-85
Summary Address configuration 56-96
Summary Address tab 56-95
Virtual Link configuration 56-90
Virtual Link MD5 configuration 56-91
Virtual Link tab 56-89
OSPFv3
LSAs 56-101
OSPFv3 routing
PIX/ASA/FWSM
advanced settings 56-104
Area/Area networks 56-108
Area Range 56-110
Area tab 56-108
Interface configuration 56-114
Interface tab 56-114
policy 56-100
Process tab 56-103
Redistribution rule 56-112
static neighbor 56-118
Summary Prefix configuration 56-113
Virtual Link configuration 56-111
OS version mismatches
handling 8-13
other settings
configuring for SSL VPN (ASA) 31-51
out-of-band changes
avoiding 8-47
detecting and analyzing 8-45
understanding 8-12
overrides
allowing overrides 6-18
creating for multiple devices 6-19
creating for single device 6-19
deleting 6-21
managing 6-17
understanding 6-18
overview
activities 1-20
device monitoring 1-7
IPv6 support 1-8
policies 1-20
ticketing 1-20
user permissions 1-11
workflow 1-20

P

P2P applications
match conditions for zone-based firewalls 21-21
P2P policy map objects
creating 21-16
match conditions and actions 21-35
packageMonitorInterval 44-6
packet capture 72-30

Packet Capture Wizard command 1-35
packet tracer 72-23
Pair dialog box 45-11
PAM
zone-based firewall
configuring 21-69
parameter maps
understanding 6-78
partial_backup.pl command 10-30
partial mesh topologies 25-5
participation, network
configuring 42-7
data collected 42-3
requirements and limitations 42-4
understanding 42-3
understanding global correlation 42-1
understanding reputation 42-2
passive OS fingerprinting on IPS sensors
configuring 40-21
understanding 40-19
Password Requirements policy, IPS platform 36-20
passwords
admin, changing 10-24
configuring IPS requirements 36-20
configuring IPS user account 36-18
discovery and deployment of IPS 36-17
managing IPS requirements 36-15
understanding managed and unmanaged IPS passwords 36-16
Paste command 1-31, 12-9
PAT
pools 24-41
PDF file
export HPM data as 71-31
PDM
device manager 72-15
Peers page 25-34
performance settings
configuring for SSL VPN (ASA) 31-52
performance tuning 44-6
permanent virtual connections (PVC)
Define Mapping dialog box 62-65
PVC Advanced Settings dialog box 62-66
PVC dialog box 62-56
PVC Policy page 62-55
permanent virtual connections (PVCs)
defining ATM PVCs 62-51
defining OAM management 62-54
on Cisco IOS routers 62-47
understanding
ATM management protocols 62-49
ATM service classes 62-48
ILMI 62-50
Operation, Administration, and Maintenance (OAM) 62-51
virtual paths and channels 62-47
per-session NAT rules 24-46
Add/Edit Per-Session NAT rules dialog boxes 24-47
PIM
configuring on firewall devices 55-11
PIX/ASA/FWSM
Bidirectional Neighbor Filter 55-14
Bidirectional Neighbor Filter tab 55-13
enable 55-1
Multicast Group, add/edit 55-19, 55-21
Multicast Group rule 55-17
Neighbor Filter 55-13
Neighbor Filter tab 55-12
page 55-11
PIM Protocol dialog box 55-12
Protocol tab 55-11
Rendezvous Point, add/edit 55-16
Rendezvous Points tab 55-15
Request Filter tab 55-18, 55-20
Route Tree tab 55-17
ping 72-26
Ping, TraceRoute and NSLookup command 1-35
PIX

PDM 72-15
PIX/ASA
boot image/configuration 48-10
add/edit 48-12
failover 50-17
settings 50-21
interfaces
Advanced tab 46-41
IP Type 46-58
MAC address 46-60
PPPoE Users 46-71
redundant 46-8
subinterfaces 46-7, 46-15
VPDN groups 46-72
mount point
add/edit 48-19, 48-20
mount point configuration 48-18
security contexts
allocate interfaces 59-12
configuration 59-9
viewing allocated interfaces 59-12
PIX/ASA/FWSM
AAA 48-5
Authentication tab 48-5
about AAA 48-1
bridging 47-1
clock settings 48-14
configuring banners 48-9
configuring CLI prompt 48-12
credentials 48-17
Device Access
Server Access 52-1, 53-1
device administration policies 48-1
Failover
bootstrap configuration 50-26
interface MAC address 50-23
failover
active/active 50-3
interface configuration 50-23
security context 50-26
understanding 50-1
interfaces
add/edit 46-31
Advanced settings 46-68
configuring 46-3
contexts 46-5
General tab 46-33
managing 46-26
operating modes 46-5
understanding 46-3
security contexts
about 59-1
Server Access
AUS, add/edit server 52-3, 53-2, 53-3, 53-4
AUS page 52-1
DDNS interface rule 52-19
DDNS page 52-18
DDNS update methods 52-19
DDNS update methods, add/edit 52-20
DHCP Relay, add/edit agent 52-6
DHCP Relay, add/edit server 52-7
DHCP Relay page 52-5
DHCP Server, add/edit 52-12
DHCP Server, advanced configuration 52-13
DHCP Server, options 52-13
DHCP Server page 52-10
DHCPv6 Relay, add/edit agent 52-9
DHCPv6 Relay, add/edit server 52-9
DHCPv6 Relay page 52-7
DNS page 52-14
DNS server, add 52-17
DNS server group 52-16
NTP page 52-21
NTP server configuration 52-21
SMTP page 52-22
TFTP server page 52-23
stateful
stateful 50-4

PIX/ASA/FWSM Platform
AAA
Accounting tab 48-8
Authorization tab 48-7
anti-spoofing 57-2
ARP configuration 47-5
ARP Inspection 47-5
enable/disable 47-6
ARP Table 47-3
configuring DHCP servers 52-10
configuring multicast routing 55-1
configuring routing 56-1
Device Access 49-1
console timeout 49-1
host name 51-1
HTTP configuration 49-3
HTTP page 49-2
ICMP rules 49-4
ICMP rules, add/edit 49-5
Management Access interface 49-6
Secure Shell, add/edit host 49-8
Secure Shell (SSH) 49-7, 49-8
SNMP host access 49-22
SNMP page 49-17
SNMP Trap configuration 49-19
Telnet configuration 49-29
Telnet page 49-29
user accounts 51-7
user accounts, add/edit 51-7
failover 50-10
failover configuration 50-1
failover configuration basics 50-5
floodguard 57-2
identity-aware IPS, QoS, and Connection Rules 13-21
IPS, QoS, and Connection Rules
wizard 58-6, 58-8
logging 54-1
email notifications 54-8
email recipients 54-8
embedded event manager 54-3
embedded event manager, add/edit action configuration 54-7
embedded event manager, add/edit applet 54-5
embedded event manager, add/edit syslog configuration 54-7
event lists 54-9
event lists, add/edit 54-10
filters 54-12
filters, editing 54-13
levels 54-24
message classes and IDs 54-9
message editing 54-25
message limits 54-18
message limits, add/edit 54-18
NetFlow 54-1
NetFlow, add/edit collector 54-2
rate limits, add/edit 54-19
server 54-21
set-up 54-15
syslog class 54-11
syslog message ID 54-11
syslog servers 54-27
syslog servers, add/edit 54-28
MAC Address
add/edit 47-8
MAC Address Table 47-8
MAC learning 47-9
enable/disable 47-9
Management IP address 47-10
multicast
Enable PIM and IGMP 55-1
group, add/edit 55-19, 55-21
IGMP Access Group parameters 55-5
IGMP Access Group tab 55-5
IGMP Join Group parameters 55-7
IGMP Join Group tab 55-7
IGMP page 55-2
IGMP parameters 55-4

IGMP Protocol tab 55-3
IGMP Static Group parameters 55-6
IGMP Static Group tab 55-6
MBoundary configuration 55-9
MBoundary interface configuration 55-10
MRoute configuration 55-8
Multicast Boundary Filter page 55-9
Multicast Group rule 55-17
Multicast Routes page 55-8
PIM Bidirectional Neighbor Filter 55-14
PIM Bidirectional Neighbor Filter tab 55-13
PIM Neighbor Filter 55-13
PIM Neighbor Filter tab 55-12
PIM page 55-11
PIM Protocol dialog box 55-12
PIM Protocol tab 55-11
PIM Rendezvous Point, add/edit 55-16
PIM Rendezvous Points tab 55-15
PIM Request Filter tab 55-18, 55-20
PIM Route Tree tab 55-17
NAT policies 24-18
Address Pools dialog box 24-19
Address Pools page 24-18
Advanced NAT Options dialog box 24-29
Dynamic Rules dialog box 24-23
Dynamic Rules tab 24-22
General tab 24-31
Policy Dynamic Rules dialog box 24-25
Policy Dynamic Rules tab 24-24
Select Address Pool 24-24
Static Rules dialog box 24-27
Static Rules tab 24-26
Translation Exemptions (NAT 0 ACL) dialog box 24-21
Translation Exemptions (NAT 0 ACL) tab 24-20
Translation Options page 24-16, 24-17
Translation Rules page 24-19
policy configuration 46-1
priority queues 58-4
priority queues configuration 58-4
routing
BGP 56-2, 56-3
BGP - General tab 56-5
BGP - IPv4 Family - Aggregate Address configuration 56-9, 56-22
BGP - IPv4 Family - Filter configuration 56-10
BGP - IPv4 Family - General tab 56-7, 56-21
BGP - IPv4 Family - Neighbor configuration 56-11, 56-24
BGP - IPv4 Family - Network configuration 56-17, 56-29
BGP - IPv4 Family - Redistribution configuration 56-18, 56-30
BGP - IPv4 Family - Route Injection configuration 56-19, 56-31
BGP - IPv4 Family tab 56-6, 56-20
EIGRP 56-32
EIGRP - advanced settings 56-34
EIGRP - Filter Rule configuration 56-40
EIGRP - Filter Rules tab 56-39
EIGRP - Interface configuration 56-48
EIGRP - Interfaces tab 56-47
EIGRP - neighbor configuration 56-42
EIGRP - Neighbors tab 56-41
EIGRP - redistribution configuration 56-44
EIGRP - Redistribution tab 56-42
EIGRP - Setup tab 56-36
EIGRP - Summary Address configuration 56-46
EIGRP - Summary Address tab 56-45
IPv6 Static Route configuration 56-131
IPv6 Static Route page 56-131
No Proxy ARP 56-1
OSPF 56-75
OSPF - advanced settings 56-77
OSPF - Area/Area networks 56-82
OSPF - Area Range 56-84
OSPF - Area tab 56-81
OSPF - Filtering configuration 56-93
OSPF - Filtering tab 56-92

OSPF - Filter Rule configuration 56-94
OSPF - Filter Rule tab 56-94
OSPF - General tab 56-76
OSPF - Interface configuration 56-98
OSPF - Interface tab 56-96
OSPF - Neighbors tab 56-85
OSPF - Range tab 56-84
OSPF - Redistribution rule 56-87
OSPF - Redistribution tab 56-86
OSPF - static neighbor 56-85
OSPF - Summary Address configuration 56-96
OSPF - Summary Address tab 56-95
OSPFv3 56-100
OSPFv3 - advanced settings 56-104
OSPFv3 - Area/Area networks 56-108
OSPFv3 - Area Range 56-110
OSPFv3 - Area tab 56-108
OSPFv3 - Interface configuration 56-114
OSPFv3 - Interface tab 56-114
OSPFv3 - Process tab 56-103
OSPFv3 - Redistribution rule 56-112
OSPFv3 - static neighbor 56-118
OSPFv3 - Summary Prefix configuration 56-113
OSPFv3 - Virtual Link configuration 56-111
OSPF - Virtual Link configuration 56-90
OSPF - Virtual Link MD5 configuration 56-91
OSPF - Virtual Link tab 56-89
RIP (PIX/ASA 6.3–7.1, FWSM) 56-120
RIP (PIX/ASA 6.3–7.1, FWSM) configuration 56-121
RIP (PIX/ASA 7.2+) 56-122
RIP (PIX/ASA 7.2+) Filtering 56-126
RIP (PIX/ASA 7.2+) Filtering configuration 56-127
RIP (PIX/ASA 7.2+) Interface 56-127
RIP (PIX/ASA 7.2+) Interface configuration 56-128
RIP (PIX/ASA 7.2+) Redistribution 56-125
RIP (PIX/ASA 7.2+) Redistribution configuration 56-125
RIP (PIX/ASA 7.2+) Setup 56-123
RIP page 56-119
static null 0 routing 56-128
Static Route configuration 56-130
Static Route page 56-128, 56-130
security contexts
managing 59-7
security group aware IPS, QoS, and Connection Rules 14-17
security policies 57-1
General configuration 57-3
General page 57-1
timeouts 57-4
service policy
wizard 58-6
Service Policy Rules 58-5
service policy rules 58-1
SNMP configuration 49-14
SNMP Version 3 49-15
traffic class 58-7
Unicast Reverse Path Forwarding 57-2
user preferences 60-1
Deployment page 60-1
Transactional Commit page 60-2
PIX/ASA/FWSM Platform policies
bridging 47-1
configuring fragment settings 57-2
configuring NAT 24-15
transparent mode 24-16
PIX 6.3
Failover
interface configuration 50-12
failover 50-10
interface configuration
IP Type 46-30
interfaces
add/edit 46-28
PIX 7.x
Failover

Add Failover Group 50-25
PIX devices
AAA support 6-28
about 46-1
monitoring service level agreements 51-8
remote access VPNs
IPsec proposals 31-41
user group policies for PIX 6.3 33-13
selecting policy types to manage 5-11
PIX Firewall
setting up AUS or CNS 2-8
setting up SSL (HTTPS) 2-3
PIX Firewalls
configuring transparent firewall rules 23-1
rollback, commands to recover from failover misconfiguration 8-68
rollback command conflicts 8-67
rollback restrictions for failover devices 8-65
rollback restrictions for multiple context mode 8-64
PIX firewalls
access controls
access list compilation 16-28
object group search 16-27
adding SSL thumbprints manually 9-5
FlexConfig object samples 7-23
packet capture, using 72-30
packet tracer, using 72-23
SSL certificate configuration 11-22
PKI (Public Key Infrastructure) policies
CA server authentication methods 26-51
defining multiple CA servers 26-55
enrollment requirements 26-52
understanding 26-51
using TFTP 26-53
PKI enrollment
prerequisites using TFTP 26-53
requirements 26-52
PKI enrollment objects
defining CA server properties 26-60
defining certificate attributes 26-66
defining enrollment parameters 26-63
defining trusted CA hierarchy 26-67
properties 26-58
plug ins
configuring browser 31-60
Point-to-Point Protocol (PPP)
defining connections 62-72
defining multilink PPP bundles 62-75
on Cisco IOS routers 62-71
understanding multilink PPP (MLP) 62-71
Point-to-Point protocol (PPP)
PPP/MLP Policy page 62-76
PPP dialog box 62-77
point-to-point topologies
description 25-3
policies
adding local rules to shared policies 5-45
assigning shared policies 5-44
basic concepts
inheritance vs. assignment 5-6
local vs. shared 5-3
managing 5-31
overview 5-1
rule inheritance 5-4
service vs. platform-specific 5-2
settings-based vs. rule-based 5-2
shared policies in Device view or Site-to-Site VPN Manager 5-37
signature inheritance 39-3
status icons 5-30
cloning shared policies 5-47
configuring IKE and IPsec for VPNs 26-1
copying between devices 5-33
creating shared 5-54
deleting shared 5-56
Device view
configuring local policies 5-31
managing 5-30

modifying assignments 5-49
modifying shared policies 5-49
discovering 5-12
discovering on existing devices 5-15
exporting 10-12
exporting with device inventory 10-6
FlexConfigs
adding objects 7-35
changing object order 7-35
changing variable values 7-35
configuring 7-25
configuring AAA for administrative introducers 63-84
editing 7-35
FlexConfig Policy page 7-36
previewing CLI 7-35
removing objects 7-35
understanding 7-2
importing 10-13
inheriting rules 5-47
locking 5-8
managing 5-1
object selectors 6-2
overview 1-20
performing basic policy management in Map view 35-22
PKI (Public Key Infrastructure) 26-51
policy banner 5-38
policy discovery FAQ 5-27
policy management and objects 5-7
Policy view
managing 5-50
modifying assignments 5-54
preshared keys 26-47
renaming 5-48
router platform policies 61-1
selecting policies to manage 5-11
sharing local 5-41
sharing multiple local policies 5-42
sharing with PRSM 72-22
Site-to-Site VPN Manager
managing 5-30
modifying assignments 5-49
site-to-site VPNs 25-8
specifying interfaces 6-76
specifying IP addresses 6-87
synchronizing among Security Manager servers 10-5
unassigning 5-36
unsharing 5-43
using global search to find specific policies 1-42
viewing discovery task status 5-22
VPN defaults 11-74
policy assignments
modifying in Device view 5-49
modifying in Policy view 5-54
modifying in Site-to-Site VPN Manager 5-49
overview 1-20
policy bundles
cloning 5-58
creating 5-57
managing 5-57
renaming 5-58, 5-59
Policy Bundle view
cloning policy bundles 5-58
creating policy bundles 5-57
renaming policy bundles 5-58, 5-59
Policy Bundle View command 1-32
policy discovery
AAA commands not displayed in AAA policy 5-29
ACL naming conventions 12-5
ACLs 5-14
Catalyst devices 5-13
Catalyst switches and 7600 Series routers 68-1
Cisco IOS routers 5-13, 61-3
frequently asked questions 5-27
IPS devices 5-14
network masks 6-82
object groups 5-14

on existing devices 5-15
overview 1-20
policy objects 5-14
preserving ACL names 12-4
resolving ACL naming conflicts 12-7
security contexts 5-13
understanding 5-12
viewing task status 5-22
VPNs 5-12
web VPN restrictions 3-8
Policy Discovery Status command 1-34
Policy Discovery Status page 5-25
Policy Dynamic Translation Rule
PIX/ASA/FWSM 24-24
add/edit 24-25
policy list objects
properties 56-143
policy management
Settings page 11-64
Policy Management page 11-64
policy maps
understanding 6-78
Policy menu
command reference 1-32
Policy Object Manager
field reference 6-4
shortcut menu 6-8
undocking and docking the window 6-8
Policy Object Manager window
creating overrides 6-19
deleting overrides 6-21
Policy Object Overrides window 6-20
policy objects
AAA server
HTTP-FORM settings 6-44
Kerberos settings 6-39
LDAP settings 6-40
NT settings 6-43
RADIUS settings 6-35
SDI settings 6-43
TACACS+ settings 6-38
AAA server groups
attributes 6-49
creating 6-48
default server groups on IOS devices 6-31
predefined authentication groups 6-30
understanding 6-27
AAA servers
creating 6-32
supported additional types for ASA/PIX/FWSM 6-28
supported types 6-28
understanding 6-27
access control lists
creating 6-53
extended objects 6-54
standard objects 6-56, 6-58
web objects 6-57
ASA group policies
client configuration settings 34-6
client firewall attributes 34-7
connection settings 34-33
DNS/WINS settings 34-29, 34-30
hardware client attributes 34-9
IPSec settings 34-10
split tunneling settings 34-31
SSL VPN clientless settings 34-12
SSL VPN full client settings 34-19
SSL VPN settings 34-25
technology settings 34-1
AS paths
properties 56-151
basic procedures 6-9
categories, using 6-13
changes in Security Manager 4.4 1-10
Cisco Secure Desktop configuration
creating 33-18
class map

creating for inspection rules 17-22
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
cloning (duplicating) 6-14
community lists
properties 56-153
configuring for ASA routing policies 56-132
configuring for remote access VPN 34-1
connection with policy management 5-7
creating 6-9
credentials
attributes 28-9
DCE/RPC policy map
properties 17-29
deleting 6-16
DNS policy map
properties 17-32
editing 6-12
ESMTP policy map
properties 17-39
exporting 6-23
file objects
attributes 34-37
selecting 34-39
FlexConfig
creating text objects 7-32
properties 7-30
property selector 7-34
undefined variables 7-33
FlexConfigs
adding to policies 7-35
changing order in policies 7-35
changing variable values 7-35
configuring 7-25
configuring AAA for administrative introducers 63-84
creating 7-28
previewing CLI 7-35
removing from policies 7-35
system variables 7-7
understanding 7-2
variables 7-5, 7-6
FTP policy map
properties 17-42
generating usage reports 6-15
GTP policy map
properties 17-45
H.323 (ASA/PIX/FWSM) policy map
properties 17-51
HTTP (ASA7.1.x/PIX7.1.x/FWSM3.x/IOS) policy map
properties 17-56
HTTP (ASA7.2+/PIX7.2+) policy map
properties 17-64
identity user group
creating 13-19
selecting 13-21
user identity acquisition 13-2
IKE proposals
v1 properties 26-10
v2 properties 26-14
IM (ASA7.2+/PIX7.2+) policy map
properties 17-70
IM (IOS) policy map
properties 17-73
importing 6-23
Inspect parameter map
properties 21-31
interface roles
creating 6-74
understanding 6-73
IP Options policy map
properties 17-75
IPsec Pass Through policy map
properties 17-80
IPSec transform sets
attributes 26-27

understanding 26-20
IPv6 policy map
properties 17-77, 17-91
LDAP attribute map objects
attributes 6-46
Local Web Filter parameter map
properties 21-38
managing 6-1
maps
understanding 6-78
N2H2 parameter map
properties 21-39
NetBIOS policy map
properties 17-81
network/host
optimizing when deploying firewall rules 12-35
understanding 6-80
using in Event Viewer filters 69-67
network/host objects
naming when provisioned as object groups 6-107
networks/hosts
creating 6-82
unspecified value objects 6-86
object selectors 6-2
overrides 3-52
allowing 6-18
creating for multiple devices 6-19
creating for single device 6-19
deleting 6-21
managing 6-17
understanding 6-18
overview 1-20
parameter map
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
PKI enrollments
defining CA server properties 26-60
defining certificate attributes 26-66
defining enrollment parameters 26-63
defining trusted CA hierarchy 26-67
properties 26-58
policy discovery 5-14
policy lists
properties 56-143
policy map
creating for inspection rules 17-22
creating for zone-based firewall content filtering 21-36
creating for zone-based firewall inspection 21-16
pools
understanding 6-92
port forwarding lists
properties 34-40
port list objects
naming when provisioned as object groups 6-107
port lists
creating 6-100
properties 6-102
prefix lists
properties 56-146, 56-148
Protocol Info parameter map
properties 21-33
provisioning as object groups 6-106
regular expression group policy map
properties 17-108
regular expression objects
metacharacters 17-109
regular expression policy map
properties 17-108
route maps 56-136
creating 56-132
understanding 56-132
security group
creating 14-14
selecting for policies 6-2
service objects
naming when provisioned as object groups 6-107

provisioning as object groups 6-108
services
creating 6-100
Settings page 11-66
sharing with PRSM 72-22
single sign-on server
properties 34-42
SIP (ASA/PIX/FWSM) policy map
properties 17-83, 17-93, 17-102, 17-103
Skinny policy map
properties 17-87
SLA monitors
attributes 51-10
configuring 51-9
understanding 51-8
SNMP policy map
properties 17-90
SSL VPN bookmark
configuring 31-82
post URL method and macro substitutions 31-84
SSL VPN Customization
configuring 31-78
creating custom Logon page 31-82
localizing 31-80
SSL VPN gateway
properties 34-64
SSL VPN smart tunnel auto sign-on lists
attributes 34-71
SSL VPN smart tunnel lists
attributes 34-66, 34-69
configuring 31-85
TCP Map policy map
properties 58-22
text
creating 7-32
time ranges
attributes for recurring ranges 6-72
configuring 6-71
traffic flow
default inspection traffic 58-20
properties 58-18
Trend parameter map
properties 21-42
TrustSec security group
selecting 14-16
URLF Glob parameter map
properties 21-45
URLF Glob parameter maps
metacharacters 21-46
URL Filter parameter map
properties 21-43
user groups
advanced PIX 6.3 settings 34-82
browser proxy settings 34-87
clientless settings 34-83
client VPN software update (IOS) settings 34-81
DNS/WINS settings 34-77
general settings 34-75
IOS client settings 34-78
IOS Xauth settings 34-80
split tunneling settings (Easy VPN/remote access IPSec VPN) 34-77
SSL VPN connection settings 34-88
SSL VPN full tunnel settings 34-84
SSL VPN split tunneling settings 34-86
technology settings 34-73
thin client settings 34-84
viewing details 6-14
Web Filter policy map
properties 21-47
Websense parameter map
properties 21-39
WINS server lists
attributes 34-90
creating 31-88
Policy Objects command 1-34
policy objects interface
Interface Role dialog box 6-75

SSL VPN Bookmark Entry dialog box 34-45
SSL VPN bookmarks
Add or Edit Bookmarks dialog boxes 34-44
Post Parameters dialog box 34-48
Policy Objects page 11-66
policy query
example report 12-34
generating reports 12-28
interpreting report results 12-32
Querying Device or Policy dialog box 12-29
Policy Query Results dialog box 12-32
Policy view
Assignments tab 5-54
creating shared policies 5-54
deleting shared policies 5-56
filtering shared policy selector 1-45
modifying assignments 5-54
overview 1-16
selectors 5-52
Shared Policy selector options 5-53
understanding 5-50
Policy View command 1-32
pool objects
understanding 6-92
POP3
configuring for inspection rules 17-21
POP3 class map objects
creating 21-16
match criteria 21-25
POP3 policy map objects
creating 21-16
match conditions and actions 21-35
port application mapping
see PAM 21-69
port forwarding list objects
properties 34-40
port list objects
creating 6-100
naming when provisioned as object groups 6-107
properties 6-102
ports
ASA 5505
configure 46-61
Posture ACL dialog box 36-29
PPP dialog box
MLP tab 62-80
PPP tab 62-78
PPPoE Users 46-71
preferences, user
PIX/ASA/FWSM 60-1
Deployment page 60-1
Transactional Commit page 60-2
prefix list objects
properties 56-146, 56-148
pre-provisioning devices 3-26
preshared keys
aggressive mode negotiation 26-48
compared to certificates 26-8
configuring policies for IKEv1 site-to-site VPNs 26-48
FQDN (fully qualified domain name) negotiation 26-48
main mode address negotiation 26-47
understanding 26-47
Preview Configuration command 1-35
Prime Security Manager
see PRSM 72-20
Prime Security Manager command 1-37
print
Report Manager reports 70-27
Print command 1-31
priority queues
PIX/ASA/FWSM
configuration 58-4
page 58-4
Product Authorization Key (PAK) 10-16
productivity categories for Trend class maps 21-20
prompt

configuring on firewall devices 48-12
properties
changes with policy effects 3-54
changing critical device 3-52
image version changes with no policy effects 3-53
understanding device 3-6
viewing or changing device 3-40
Property Selector dialog box 7-34
protected networks
defining in GET VPN topologies 25-60
defining in VPN topologies 25-34
Protected Networks tab 25-46
Protocol Independent Multicast 55-11
Protocol Info parameter map objects
properties 21-33
Protocol Info Parameters map object
creating 21-16
Protocol Map dialog box 41-12
protocols
selecting for inspection 17-3
Protocol tab
IGMP 55-3
proxies
defining HTTP/HTTPS for SSL VPN (ASA) 31-57
proxy ARP
enabling on IOS routers 62-19
proxy bypass rules
defining HTTP/HTTPS for SSL VPN (ASA) 31-57
proxy server
configuring HTTP for IPS global correlation 36-24
PRSM
sharing
devices 72-22
policy objects 72-22
starting from Security Manager 72-20
public key infrastructure (PKI) policies
compared to certificates 26-8
configuring for remote access VPNs 26-56
configuring for site-to-site VPNs 26-54
PVC Advanced Settings dialog box
OAM-PVC tab 62-69
OAM tab 62-67
PVC dialog box
Protocol tab 62-64
QoS tab 62-61
Settings tab 62-58
PVC policies
unable to deploy 9-15

Q

QoS
MPC rule wizard
tab 58-8
PIX/ASA/FWSM
identity-aware rules 13-21
rules 58-5
QoS Class dialog box 66-23
Edit ACLs dialog box 66-25
Marking tab 66-26
Matching tab 66-24
Policing tab 66-29
Queuing and Congestion Avoidance tab 66-27
Shaping tab 66-31
QoS queuing
default class 66-6
defining for classes 66-16
tail drop vs. WRED 66-4
understanding 66-4
understanding LLQ 66-5
quality of service (QoS)
CEF requirements 66-2
defining on control plane 66-12
defining on interfaces 66-10
defining policies 66-10
on Cisco IOS routers 66-1
QoS Class dialog box 66-23
QoS Policy dialog box 66-21

Quality of Service Policy page 66-19
understanding
Control Plane Policing 66-9
default class queuing 66-6
low-latency queuing 66-5
marking parameters 66-3
matching parameters 66-2
policing parameters 66-6
queuing parameters 66-4
shaping parameters 66-6
tail drop and WRED 66-4
token-bucket mechanism 66-8
quality of service (QoS) classes
defining marking parameters 66-15
defining matching parameters 66-13
defining policing parameters 66-17
defining queuing parameters 66-16
defining shaping parameters 66-18
query
CS-MARS
access rule events 72-42
IPS signature events 72-44
looking up policies based on related events 72-45
overview 72-41
troubleshooting 72-40
Event Viewer
access rule events 69-56
IPS signature events 69-57
looking up policies based on related events 69-54
overview 69-55
Querying Device or Policy dialog box 12-29
quick filter
searching for events 69-47
Quick Launch for IP Intelligence 72-35

R

RADIUS
description 6-28
settings in AAA server objects 6-35
RAM
Image Manager 73-17
rate limiting, IPS 43-4
Real-time Log Viewer 72-18
recovery
event data store 69-36
Recurring Ranges dialog box 6-72
Redeploy a Job dialog box 8-53
Redeploying Licenses dialog box 11-60
rediscovering
remote access VPNs 30-12
rediscovering site-to-site VPNs 25-27
Rediscover VPN Policies wizard 25-27
redundant interfaces 46-8
red X in device selector, troubleshooting 9-9
Refresh Map command 1-33
regular expression group objects
properties 17-108
regular expression objects
metacharacters 17-109
properties 17-108
regular IPsec
mandatory and optional policies 25-6
supported platforms 25-9
supported platforms for remote access VPNs 30-8
Reject Activity command 1-36
Reject Activity dialog box 4-21
Reject Deployment Job dialog box 8-20, 8-39
remote access
user
logging off 71-30
remote access VPN
system variables 7-19
Remote Access VPN Configuration wizard
IPsec VPN
Defaults page 30-31
IPsec Settings page (ASA) 30-30
IPsec VPN Connection Profile page (ASA) 30-28

User Groups page 30-36
IPsec VPNs
creating on ASA/PIX 7.0+ 30-25
creating on IOS/PIX 6.3+ 30-36
SSL VPN
Access page (ASA) 30-16
Connection Profile page (ASA) 30-17
Gateway and Context Page (IOS) 30-33
Portal Page Customization Page (IOS) 30-35
SSL VPNs
creating on ASA devices 30-14
creating on IOS devices 30-32
using 30-13
remote access VPNs
ASA devices
configuring bookmarks 31-82
configuring portal appearance 31-78
configuring WINS servers for file system access 31-88
customizing 31-77
post URL method and macro substitutions in bookmarks 31-84
smart tunnels 31-85
configuring using wizard 30-13
device support 30-8
discovering 30-12
IOS devices
configuring bookmarks 31-82
configuring WINS servers for file system access 31-88
IPsec 31-36
access policies for IKEv2 (ASA), configuring 31-49
access policies for IKEv2 (ASA), reference 31-45
access policies for IKEv2 (ASA), understanding 31-44
certificate to connection profile map policy (IKEv1) 31-36
certificate to connection profile map rules (IKEv1) 31-37
cluster load balancing 31-5
configuring IKE and IPsec policies 26-1
connection profiles 31-7
connection profiles (ASA, PIX 7+) 31-8
creating on ASA/PIX 7.0+ 30-25
creating on IOS/PIX 6.3+ 30-36
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
Dynamic VTI/VRF Aware IPsec settings 33-7
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
high availability policies 33-11
IKE proposals 26-9
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
NAT settings 26-42
policy overview 30-9
policy overview (ASA, PIX 7.0+) 31-2
policy overview (IOS, PIX 6.3) 33-2
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9
understanding 30-2
understanding IKE 26-5
understanding NAT settings 26-41
user group policies for IOS, PIX 6.3 33-13
VPNSM, VPN SPA, VSPA settings 33-6
IPsec proposals
attributes for ASA and PIX 7.0+ devices 31-41
attributes for IOS and PIX 6.3 devices 33-4
configuring for ASA and PIX 7.0+ devices 31-40
configuring for IOS and PIX 6.3 devices 33-3
managing 30-1
managing (ASA, PIX 7.0+) 31-1
managing (IOS, PIX 6.3) 33-1
rediscovering 30-12
SSL 31-43

access modes 30-4
access policies (ASA), configuring 31-49
access policies (ASA), reference 31-45
access policies (ASA), understanding 31-44
advanced settings (ASA) 31-72
AnyConnect client image settings (ASA) 31-65
AnyConnect client settings (ASA) 31-62, 31-64
AnyConnect custom attributes(ASA) 31-70, 31-71
browser plug-ins (ASA) 31-60
cluster load balancing 31-5
configuring HTTP/HTTPS proxies and proxy bypass(ASA) 31-57
connection profiles 31-7
connection profiles (ASA) 31-8
content rewrite rules (ASA) 31-53
Context Editor dialog box (IOS) 33-15, 33-16
creating on ASA 30-14
creating on IOS devices 30-32
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
encoding rules (ASA) 31-55
example 30-3
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
Kerberos Constrained Delegation (KCD on ASA) 31-66, 31-69
limitations 30-8
managing support files 30-5
NAT settings 26-42
other settings (ASA) 31-51
performance settings (ASA) 31-52
policies (IOS) 33-14
policy overview 30-9
policy overview (ASA, PIX 7.0+) 31-2
policy overview (IOS, PIX 6.3) 33-2
prerequisites 30-7
proxy bypass rules (ASA) 31-59
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9
server certificate verification (ASA) 31-30, 31-32, 31-73
shared license (ASA) 31-74
shared license clients (ASA) 31-76
shared license servers (ASA) 31-77
understanding 30-2
understanding NAT settings 26-41
wizard 30-13
understanding 30-1
Remote Detection Indication (RDI) cells 62-51
Rename Policy Bundle dialog box 5-58, 5-59
Rename Policy command 1-32
Rename Policy dialog box 5-48
Rendezvous Point
PIX/ASA/FWSM
add/edit 55-16
Rendezvous Points
PIM 55-15
Report Manager
arranging window 70-30
closing 70-32
configuring default settings 70-29
configuring devices to provide reports 70-3
configuring Event Manager service 69-30
configuring schedules 70-34
creating custom reports 70-21
deleting another user's custom reports 70-32
deleting reports 70-32
deleting schedules 70-36
disabling schedules 70-36
drill-down reports 70-26
editing report settings 70-22
enabling schedules 70-36

examples of analysis
monitoring botnet activity 69-64
exporting reports 70-28
generated report pane and toolbar 70-12
generating reports 70-20
managing custom reports 70-32
opening reports 70-20
overview 70-1, 70-6
printing reports 70-27
renaming reports 70-31
report list 70-9
report settings 70-10
saving reports 70-31
scheduling reports 70-33
settings page 11-38
troubleshooting 70-36
understanding 70-1
understanding access control 70-5
understanding data aggregation 70-4
understanding predefined reports
firewall summary botnet reports 70-15
firewall traffic reports 70-14
general IPS reports 70-19
general VPN reports 70-16
IPS top reports 70-17
overview 70-13
VPN top reports 70-16
using 70-19
viewing schedule results 70-35
viewing schedules 70-33
Report Manager command 1-38
Report Manager in Dashboard 72-2
reports
arranging windows 70-30
closing 70-32
configuring default settings for reports 70-29
configuring devices for Report Manager reporting 70-3
configuring schedules 70-34
creating custom 70-21
deleting 70-32
deleting another user's in Report Manager 70-32
deleting schedules 70-36
deployment status 10-30
disabling schedules 70-36
discovery status 10-30
drilling down 70-26
editing settings 70-22
enabling schedules 70-36
example policy query 12-34
exporting 70-28
generating 70-20
generating access rule analysis 16-34
generating policy query 12-28
interpreting policy query 12-32
managing 70-1
managing custom 70-32
opening 70-20
overview of available types 70-2
predefined Report Manager
firewall summary botnet reports 70-15
firewall traffic reports 70-14
general IPS reports 70-19
general VPN reports 70-16
IPS top reports 70-17
overview 70-13
VPN top reports 70-16
printing 70-27
renaming 70-31
Report Manager
generated report pane and toolbar 70-12
overview 70-6
report list 70-9
report settings 70-10
saving 70-31
scheduling in Report Manager 70-33
understanding Report Manager 70-1
understanding Report Manager access control 70-5

understanding Report Manager data aggregation 70-4
using Report Manager 70-19
viewing schedule results 70-35
viewing schedules 70-33
reputation
configuring global correlation 42-5
understanding IPS global correlation 42-2
Request Filter
PIM 55-18, 55-20
Resources
FWSM 51-3
add/edit 51-4
managing 51-2
restorebackup.pl command 10-27
restore database 10-27
Resume Deployment Schedule dialog box 8-20, 8-58
retry count
device communication 11-22
reverse route injection 26-21
RIP
PIX/ASA/FWSM 56-119
(PIX/ASA 6.3–7.1, FWSM) 56-120
(PIX/ASA 6.3–7.1, FWSM) configuration 56-121
(PIX/ASA 7.2+) 56-122
(PIX/ASA 7.2+) Filtering 56-126
(PIX/ASA 7.2+) Filtering configuration 56-127
(PIX/ASA 7.2+) Interface 56-127
(PIX/ASA 7.2+) Interface configuration 56-128
(PIX/ASA 7.2+) Redistribution 56-125
(PIX/ASA 7.2+) Redistribution configuration 56-125
(PIX/ASA 7.2+) Setup 56-123
RIP routing
Cisco IOS routers
Authentication dialog box 67-47
Authentication tab 67-46
defining interface authentication 67-43
defining setup parameters 67-42
overview 67-42
redistributing routes 67-44
Redistribution Mapping dialog box 67-49
Redistribution tab 67-48
RIP Routing Policy page 67-45
Setup tab 67-45
roles, IPS user 36-15
rollback
archived configuration files 8-70
last deployed configuration 8-69
when deploying to file 8-71
Rollback a Job dialog box 8-69
round robin allocation
PAT 24-41
routed ports
Create and Edit Interface dialog boxes-Routed Port mode 68-12
understanding 68-5
route map objects
creating 56-132
properties 56-136
understanding 56-132
Router Block Interface dialog box 43-15
Router Device dialog box 43-14
router platform interface
802.1x Policy page 64-5
AAA policy
AAA Policy page 63-6
Accounting tab 63-10
Authentication tab 63-6
Authorization tab 63-8
Command Accounting dialog box 63-13
Command Authorization dialog box 63-10
accounts and credentials policy
Accounts and Credentials Policy page 63-16
User Accounts dialog box 63-17
ADSL policy
ADSL Policy page 62-37
ADSL Settings dialog box 62-38
advanced interface settings policy

Advanced Interface Settings dialog box 62-16
Advanced Interface Settings page 62-16
BGP policy
BGP Neighbors dialog box 67-6
BGP Redistribution tab 67-7
BGP Routing Policy page 67-4
BGP Setup tab 67-5
Redistribution Mapping dialog box 67-7
bridging policy
Bridge Group dialog box 63-21
Bridging Policy page 63-21
CEF interface policy 62-26
CEF Interface Settings dialog box 62-27
Clock Policy page 63-23
console policy
AAA tab 63-44
Accounting tab 63-47
Authentication tab 63-44
Authorization tab 63-45
Command Accounting dialog box 63-61
Command Authorization dialog box 63-60
Console Policy page 63-42
Setup tab 63-42
CPU Policy page 63-26
DHCP policy
DHCP Database dialog box 63-94
DHCP Policy page 63-92
IP Pool dialog box 63-94
dialer interface policy
Dialer Physical Interface dialog box 62-33
Dialer Policy page 62-31
Dialer Profile dialog box 62-32
DNS policy
IP Host dialog box 63-76
DNS Policy page 63-76
EIGRP policy
EIGRP Routing Policy page 67-13
Interface dialog box 67-16
Interfaces tab 67-15
Redistribution Mapping dialog box 67-18
Redistribution tab 67-17
Setup dialog box 67-14
Setup tab 67-13
Hostname Policy page 63-78
HTTP policy
AAA tab 63-32
Command Authorization Override dialog box 63-34
HTTP Policy page 63-31
Setup tab 63-31
interfaces policy
Create Router Interface dialog box 62-8
Interface Auto Name Generator dialog box 62-12
Router Interfaces page 62-7
IPS interface policy
IPS Monitoring Information dialog box 62-24
IPS Module interface policy
IPS Module Interface Policy Page 62-23
logging policy
Syslog Server dialog box 65-11
logging setup policy
Logging Setup Policy page 65-7
Memory Policy page 63-79
NAC policy
Identities tab 64-18
Identity Action dialog box 64-19
Identity Profile dialog box 64-19
Interface Configuration dialog box 64-17
Interfaces tab 64-16
NAC Policy page 64-14
Setup tab 64-14
NAT policy
Dynamic Rule dialog box 24-11
Interface Specification tab 24-6
Static Rule dialog box 24-7
Static Rules tab 24-6
NetFlow policy 65-5, 65-12
NTP policy

NTP Policy page 63-98
NTP Server dialog box 63-99
OSPF policy
Area dialog box 67-37
Area tab 67-36
Interface dialog box 67-31
Max Prefix Mapping dialog box 67-41
OSPF Interface Policy page 67-30
OSPF Process Policy page 67-34
Redistribution Mapping dialog box 67-39
Redistribution tab 67-38
Setup dialog box 67-35
Setup tab 67-35
PPP/MLP policy
PPP/MLP Policy page 62-76
PPP dialog box 62-77
PVC policy
Define Mapping dialog box 62-65
PVC Advanced Settings dialog box 62-66
PVC dialog box 62-56
PVC Policy page 62-55
QoS policy
QoS Class dialog box 66-23
QoS Policy dialog box 66-21
Quality of Service Policy page 66-19
RIP policy
Authentication dialog box 67-47
Authentication tab 67-46
Redistribution Mapping dialog box 67-49
Redistribution tab 67-48
RIP Routing Policy page 67-45
Setup tab 67-45
Secure Device Provisioning Policy page 63-85
Secure Shell Policy page 63-64
SHDSL policy
Controller Auto Name Generator dialog box 62-46
SHDSL Controller dialog box 62-43
SHDSL Policy page 62-42
SNMP policy
Permission dialog box 63-70
SNMP Policy page 63-69
SNMP Traps dialog box 63-72
Trap Receiver dialog box 63-71
static routing policy
Static Routing dialog box 67-52
Static Routing Policy page 67-51
syslog servers policy
Syslog Servers Policy page 65-10
VTY policy
Command Accounting dialog box 63-61
Command Authorization dialog box 63-60
VTY Line dialog box 63-51
VTY Policy page 63-50
router platform policies
Device Admin policies
AAA 63-2
accounts and credentials 63-14
CPU settings 63-25
DHCP 63-87
DNS 63-74
host and domain names 63-77
HTTP 63-28
line access 63-35
memory settings 63-78
optional SSH settings 63-63
Secure Device Provisioning (SDP) 63-81
SNMP 63-66
time zone settings 63-22
transparent bridging 63-18
Identity policies
802.1x 64-1
Network Admission Control (NAC) 64-8
Interface policies
ADSL 62-34
advanced settings 62-13
basic settings 62-1
dialer interfaces 62-28

PPP 62-71
PVC 62-47
SHDSL 62-41
Logging policies 65-1
NAT 24-5
dynamic rules 24-10
static rules 24-6
timeouts 24-13
NetFlow policies 65-1
Network Time Protocol (NTP) 63-96
quality of service (QoS) 66-1
Routing policies
BGP routing 67-1
EIGRP routing 67-8
OSPF routing 67-19
RIP routing 67-42
static routing 67-50
routers
adding SSL thumbprints manually 9-5
CEF interface settings policies 62-25
Cisco Discovery Protocol (CDP) settings 62-18
communication requirements 2-1
configuring SSH 2-6
default transport protocol for 12.1 and 12.2 11-22
default transport protocol for 12.3 and above 11-22
deploying configurations using TMS 8-43
enabling directed broadcasts 62-20
enabling Maintenance Operation Protocol (MOP) 62-19
enabling NBAR protocol discovery 62-19
enabling proxy ARP 62-19
enabling unicast reverse path forwarding (RFP) 62-20
enabling virtual fragment reassembly (VFR) 62-19
FlexConfig object samples 7-24
generating interface names 62-4
ICMP message settings 62-18
IPS Module interface settings policies 62-22
licenses 2-10
mixing deployment methods 9-14
selecting policy types to manage 5-11
setting up SSL (HTTPS) 2-4
SSL certificate configuration 11-22
system variables 7-13
troubleshooting deployment 9-14
Route Tree
PIM 55-17
routing
PIX/ASA/FWSM
about EIGRP 56-33
about OSPF 56-75
about OSPFv3 56-101
authentication 56-75
BGP 56-2, 56-3
BGP - General tab 56-5
BGP - IPv4 Family - Aggregate Address configuration 56-9, 56-22
BGP - IPv4 Family - Filter configuration 56-10
BGP - IPv4 Family - General tab 56-7, 56-21
BGP - IPv4 Family - Neighbor configuration 56-11, 56-24
BGP - IPv4 Family - Network configuration 56-17, 56-29
BGP - IPv4 Family - Redistribution configuration 56-18, 56-30
BGP - IPv4 Family - Route Injection configuration 56-19, 56-31
BGP - IPv4 Family tab 56-6, 56-20
configuring on 56-1
configuring static routes 56-128
EIGRP 56-32
EIGRP - advanced settings 56-34
EIGRP - Filter Rule configuration 56-40
EIGRP - Filter Rules tab 56-39
EIGRP - Interface configuration 56-48
EIGRP - Interfaces tab 56-47
EIGRP - neighbor configuration 56-42
EIGRP - Neighbors tab 56-41
EIGRP - redistribution configuration 56-44
EIGRP - Redistribution tab 56-42

EIGRP - Setup tab 56-36
EIGRP - Summary Address configuration 56-46
EIGRP - Summary Address tab 56-45
IPv6 Static Route configuration 56-131
No Proxy ARP 56-1
OSPF 56-75
OSPF - advanced settings 56-77
OSPF - Area/Area networks 56-82
OSPF - Area Range 56-84
OSPF - Area tab 56-81
OSPF - Filtering configuration 56-93
OSPF - Filtering tab 56-92
OSPF - Filter Rule configuration 56-94
OSPF - Filter Rule tab 56-94
OSPF - General tab 56-76
OSPF - Interface configuration 56-98
OSPF - Interface tab 56-96
OSPF - Neighbors tab 56-85
OSPF - Range tab 56-84
OSPF - Redistribution rule 56-87
OSPF - Redistribution tab 56-86
OSPF - static neighbor 56-85
OSPF - Summary Address configuration 56-96
OSPF - Summary Address tab 56-95
OSPFv3 56-100, 56-101
OSPFv3 - advanced settings 56-104
OSPFv3 - Area/Area networks 56-108
OSPFv3 - Area Range 56-110
OSPFv3 - Area tab 56-108
OSPFv3 - Interface configuration 56-114
OSPFv3 - Interface tab 56-114
OSPFv3 - Process tab 56-103
OSPFv3 - Redistribution rule 56-112
OSPFv3 - static neighbor 56-118
OSPFv3 - Summary Prefix configuration 56-113
OSPFv3 - Virtual Link configuration 56-111
OSPF - Virtual Link configuration 56-90
OSPF - Virtual Link MD5 configuration 56-91
OSPF - Virtual Link tab 56-89
RIP (PIX/ASA 6.3–7.1, FWSM) 56-120
RIP (PIX/ASA 6.3–7.1, FWSM) configuration 56-121
RIP (PIX/ASA 7.2+) 56-122
RIP (PIX/ASA 7.2+) Filtering 56-126
RIP (PIX/ASA 7.2+) Filtering configuration 56-127
RIP (PIX/ASA 7.2+) Interface 56-127
RIP (PIX/ASA 7.2+) Interface configuration 56-128
RIP (PIX/ASA 7.2+) Redistribution 56-125
RIP (PIX/ASA 7.2+) Redistribution configuration 56-125
RIP (PIX/ASA 7.2+) Setup 56-123
RIP page 56-119
static null 0 routing 56-128
Static Route configuration 56-130
VPNs with routing processes 9-13
routing redistribution
BGP Redistribution Mapping dialog box 67-7
BGP Redistribution tab 67-7
EIGRP Redistribution Mapping dialog box 67-18
EIGRP Redistribution tab 67-17
into BGP 67-3
into EIGRP 67-12
into OSPF 67-22
into RIP 67-44
OSPF Max Prefix Mapping dialog box 67-41
OSPF Process Redistribution tab 67-38
OSPF Redistribution Mapping dialog box 67-39
RIP Redistribution Mapping dialog box 67-49
RIP Redistribution tab 67-48
RPC
configuring for inspection rules 17-21
RSA keys
generating, synchronizing for GET VPN 29-13
Rule Analysis Detail Report
generating 16-34
Rule Combiner Results dialog box 12-25
rule expiration

configuring for access rules 16-22
Rule Expiration page 11-69
rules
default 5-5
mandatory 5-5
rules tables
adding rules 12-9
columns and headings 1-49
commands, Edit menu 1-31
converting IPv4 rules 12-28
cut, copy, and paste rules 12-9
disabling rules 12-20
enabling rules 12-20
filtering 1-48
finding and replacing items 12-16
removing rules 12-9
sections 12-20
using 12-8
rule tables
moving rules 12-19
RX-Boot Mode Credentials dialog box 3-48

S

Save As command (Report Manager) 70-8
Save command 1-30
Save command (Report Manager) 70-8
Save Map As command 1-33
Save Map As dialog box 35-10
Save Map command 1-33
ScanSafe Web Security Settings 20-6
scenarios
creating FlexConfigs 7-25
SCEP (Simple Certificate Enrollment Protocol)
CA server authentication 26-51
Schedule dialog box 8-56
schedules
configuring in Report Manager 70-34
deleting in Report Manager 70-36
disabling in Report Manager 70-36
enabling in Report Manager 70-36
reports in Report Manager 70-33
viewing in Report Manager 70-33
viewing results in Report Manager 70-35
schedules, deployment
changes not deployed 8-55
creating or editing 8-55
including devices 8-8
suspending or resuming 8-58
viewing status and history 8-26
scripting language
examples
looping 7-3
looping with if/else statements 7-4
looping with two-dimensional arrays 7-3
FlexConfig objects 7-3
SDEE
subscriptions for IOS IPS 45-8
SDI
settings in AAA server objects 6-43
SDM
access rule look-up 72-19
device manager 72-16
searching for items 1-42
Secondary Interface Specific Authentication Server Groups dialog box 31-16
secure desktop manager policies
configuring 32-9
Secure Device Provisioning (SDP)
configuring AAA for administrative introducers 63-84
contents of bootstrap 63-82
defining policies 63-83
Secure Device Provisioning page 63-85
understanding
introducers 63-81
petitioners 63-81
registrars 63-81

TTI 63-81
workflow 63-82
SecureID servers (SDI)
description 6-29
Secure Shell
PIX/ASA/FWSM
add/edit SSH host 49-8
Secure Shell (SSH)
Cisco IOS routers
defining optional settings 63-63
optional settings overview 63-63
Secure Shell Policy page 63-64
PIX/ASA/FWSM 49-7, 49-8
security associations
GET VPN
using passive mode during migration 29-23
security certificate
invalid during discovery 9-7
security context
Failover page 50-26
security contexts
adding to failover group 2 50-8
admin context
overview 59-1
configuring multiple 59-3
configuring on firewall devices 59-1
deleting FWSM 59-7
discovering policies 5-13
FWSM 59-8
configuration 59-8
managing Resources 51-2
Resources 51-3
PIX/ASA
allocate interfaces 59-12
configuration 59-9
viewing allocated interfaces 59-12
PIX/ASA/FWSM
enabling multi-context mode 59-1
managing 59-7
restoring single-context mode 59-1
rollback, commands to recover from failover misconfiguration 8-68
rollback command conflicts 8-67
rollback restrictions 8-64
rollback restrictions for failover devices 8-65
showing containment 3-56
security group aware firewall policies
configuring ISE settings 11-56
security group-aware firewall policies
configuring 14-7
managing 14-1
overview 14-1
security group objects
creating 14-14
security group tagging 46-44
Security Manager
access by CS-MARS 72-37
applications overview 1-6
archiving (backing up) the event data store 69-36
backing up and restoring database 10-24
Configuration Manager interface overview 1-14
configuring administrative settings 11-1
getting started 1-1
how permissions affect what you can do 1-11
initial configuration 1-25
installing client 1-12
integrating with Security Manager 72-36
integration with CS-MARS 72-36
logging into and exiting 1-12
managing the server 10-1
overview 1-1
recovering the event data store 69-36
reports overview 70-2
server cluster
managing 10-2
overview 10-2
splitting server 10-3
synchronizing shared policies 10-5

server management and administration 10-1
using 1-14
Security Manager Administration command 1-36
Security Manager Diagnostics command 1-36
Security Manager Online command 1-38
security policies
PIX/ASA/FWSM 57-1
General configuration 57-3
General page 57-1
timeouts 57-4
security ratings for Trend class maps 21-20
see LACP 46-12
Select Address Pool
PIX/ASA/FWSM Platform 24-24
Select Interfaces dialog box 35-20
selectors
filtering items 1-45
using 1-45
selector trees
selecting items 1-45
Select Policy Object dialog box 35-18
Select VPN to Configure dialog box 35-22
self near-end crosstalk (SNEXT) 62-46
Self zone 21-5
sensors, IPS
allowed hosts 36-7
anomaly detection
configuring 41-6
configuring histograms 41-11
configuring learning accept mode 41-8
configuring signatures 41-4
configuring thresholds 41-11
detection zones 41-3
managing 41-1
modes 41-2
understanding 41-1
understanding histograms 41-9
understanding thresholds 41-9
understanding worms 41-2
when to turn off 41-4
blocking
configuring 43-7
configuring ARC 43-1
configuring blocking devices 43-14
configuring master blocking sensors 43-13
configuring never block hosts and networks 43-17
configuring router blocking interfaces 43-15
configuring user profiles 43-12
configuring VLAN blocking interfaces 43-16
general options 43-10
master blocking sensor 43-6
policy 43-8
rate limiting 43-4
router and switch blocking devices 43-4
strategies 43-3
understanding 43-1
capturing network traffic 36-2
certificates 44-10
configuration overview 36-5
configuration overview for IOS IPS 45-4
configuring AAA 36-21
configuring Analysis Engine global variables 36-30
configuring DNS servers 36-24
configuring HTTP proxy server 36-24
configuring NTP 36-23
configuring OS maps 40-21
configuring SNMP 36-8
configuring target value ratings 40-17
configuring the external product interface 36-26
configuring user accounts 36-18
deployment of passwords 36-17
deployment topology 36-4
discovery of passwords 36-17
event actions
example filter rule 69-67
filter rule attributes 40-9
filter rules 40-4, 40-7

filter rules tips 40-6
network information 40-17
overrides 40-13
overview 40-1
possible actions 40-2
process overview 40-1
settings 40-23
getting started 36-1
global correlation
configuring 42-1
configuring inspection and reputation 42-5
configuring network participation 42-7
data collected 42-3
requirements and limitations 42-4
understanding 42-1
understanding network participation 42-3
understanding reputation 42-2
interfaces
configuring 37-6
configuring bypass mode 37-12
configuring CDP mode 37-12
configuring inline interface pairs 37-13
configuring inline VLAN pairs 37-14
configuring physical 37-9
configuring VLAN groups 37-15
deploying VLAN groups 37-5
inline interface mode 37-3
inline VLAN pair mode 37-3
interfaces policy 37-6
managing interface configurations 37-1
physical interface properties 37-10
promiscuous mode 37-2
roles 37-1
sensing modes overview 37-2
understanding 37-1
viewing summary 37-8
VLAN group mode 37-4
IPS modules for ASA 58-15
licenses
automating 44-3
managing 44-1
redeploying 44-2
updating 44-1
managing 44-1
managing user accounts and passwords 36-15
monitoring
removing false positive IPS events 69-66
passive OS fingerprinting 40-19
password requirements 36-20
rebooting 44-12
signatures
adding custom 39-19
cloning 39-21
configuring 39-4
configuring settings 39-30
defining 39-1
detailed information 39-2
editing 39-14
editing Meta engine component list 39-29
editing or tuning parameters 39-23
enabling or disabling 39-14
engines 39-20
exporting 39-9
inheritance 39-3
parameters list 39-24
policy 39-4
shortcut menu 39-10
understanding 39-1
viewing update level 39-9, 39-13
traffic flow notifications 36-30
tuning recommendations 36-4
understanding managed and unmanaged passwords 36-16
understanding network sensing 36-2
understanding user roles 36-15
updates
automatically applying 44-6
checking for and downloading 44-5

configuring server 44-4
managing 44-4
manually applying 44-7
user account attributes 36-19
virtual sensors
advantages 38-3
assigning interfaces 38-4
attributes 38-7
configuring 38-1, 38-5
deleting 38-10
editing policies 38-9
identifying 38-5
inline TCP session tracking mode 38-3
Normalizer mode 38-4
renaming 38-8
restrictions 38-3
understanding 38-1
sensorupdate.properties 44-6
server
managing Security Manager 10-1
syslog
PIX/ASA/FWSM 54-21, 54-27
server, IPS update 44-4
server, Security Manager
configuring administrative settings 11-1
managing or administrating 10-1
Server Access
PIX/ASA/FWSM 52-1, 53-1
AUS, add/edit server 52-3, 53-2, 53-3, 53-4
AUS page 52-1
DDNS interface rule 52-19
DDNS page 52-18
DDNS update methods 52-19
DDNS update methods, add/edit 52-20
DHCP Relay, add/edit agent 52-6
DHCP Relay, add/edit server 52-7
DHCP Relay page 52-5
DHCP Server, add/edit 52-12
DHCP Server, advanced configuration 52-13
DHCP Server, options 52-13
DHCP Server page 52-10
DHCPv6 Relay, add/edit agent 52-9
DHCPv6 Relay, add/edit server 52-9
DHCPv6 Relay page 52-7
DNS page 52-14
DNS server, add 52-17
DNS server group 52-16
NTP page 52-21
NTP server configuration 52-21
SMTP page 52-22
TFTP server page 52-23
server cluster, Security Manager
managing 10-2
overview 10-2
splitting server 10-3
synchronizing shared policies 10-5
Server Load Balance page 27-17
server load balancing
configuring for large scale DMVPN 27-16, 27-17
server attributes in large scale DMVPN 27-17
Server Properties dialog box 3-38
Server Security page 10-2
Server Security Settings page 11-70
Service
ASA CX
Auth Proxy Configuration 58-17
PIX/ASA/FWSM
identity-aware IPS, QoS, and Connection Rules 13-21
IPS, QoS, and Connection Rules wizard 58-6, 58-8
policy wizard 58-6
priority queues 58-4
priority queues configuration 58-4
security group aware IPS, QoS, and Connection Rules 14-17
Service Policy Rules 58-5
traffic class 58-7
service, Event Manager

configuring 69-30
managing 69-30
monitoring event store disk space 69-35
monitoring status 69-31
selecting devices to monitor 69-34
starting or stopping 69-30
status icon colors 69-31
service agreement contracts 10-16
Service Contents dialog box 12-14
Service Device Provisioning (SDP)
on Cisco IOS routers 63-81
Service Module Credentials dialog box 3-19
Service Modules
Catalyst
firewalls 46-1
service objects
creating 6-100
naming when provisioned as object groups 6-107
provisioning as object groups 6-108
Services dialog box 6-103
understanding 6-100
service policy
configuring identity-aware rules 13-21
configuring security group aware rules 14-17
Service Policy (MPC) Rule Wizard 58-6
Connection Settings tab 58-8
CSC tab 58-8
CXSC tab 58-8
IPS tab 58-8
QoS tab 58-8
User Statistics tab 58-8
service policy rules
configuring on firewall devices 58-1
services
specifying 6-100
Set Linked Map dialog box 35-13
Settings
ScanSafe 20-6
settings
device communications 9-4
Settings, Event Actions policy 40-23
settings, report
editing 70-22
Settings pages
Autolink 11-3
CCO Settings 11-4
Configuration Archive 11-6
CS-MARS 11-7
CSM Mobile 11-9
Customize Desktop 11-10
Debug Options 11-11
Deployment 11-13
Device Communication 11-21
Device Groups 11-24
Discovery 11-25
Event Management 11-27, 11-35
CPU Throttling Policy 11-33
Health and Performance Monitor 11-36
Identity 11-38
Image Manager 11-41
IP Intelligence Settings 11-41
ISE 11-56
Licensing 11-57
Logs 11-62
Policy Management 11-64
Policy Objects 11-66
Report Manager 11-38
Rule Expiration 11-69
Server Security 11-70
Take Over User Session 11-71
Ticket Management 11-72
Token Management 11-73
VPN Policy Defaults 11-74
Workflow 11-75
SHA hash algorithm 26-6
Share Device Policies command 1-32
shared license clients
configuring 31-76

shared license servers
configuring 31-77
shared policies
cloning (copying) 5-47
Device view
adding local rules to selected device 5-45
assigning to selected device 5-44
modifying 5-49
modifying assignments 5-49
policy banner 5-38
sharing local 5-41
sharing multiple local policies 5-42
unsharing 5-43
working with 5-37
exporting 10-12
exporting with device inventory 10-6
importing 10-13
inheriting policies 5-47
Policy Bundle view
cloning 5-58
creating 5-57
renaming 5-58, 5-59
Policy view
creating 5-54
deleting 5-56
managing 5-50
modifying assignments 5-54
renaming 5-48
Site-to-Site VPN Manager
assigning to selected device 5-44
modifying assignments 5-49
sharing local 5-41
unsharing 5-43
working with 5-37
synchronizing among Security Manager servers 10-5
Shared Policy Assignments dialog box 5-49
Share Policies wizard 5-42
Share Policy command 1-32
Share Policy dialog box 5-41
SHDSL
Controller Auto Name Generator dialog box 62-46
defining controllers 62-41
on Cisco IOS routers 62-41
SHDSL Controller dialog box 62-43
SHDSL Policy page 62-42
shortcut menu commands
policies in Device view and Site-to-Site VPN Manager 5-40
Show Containment command 1-35
Show Devices On Map command 1-33
Show Devices on Map dialog box 35-16
Show Navigation Window command 1-33
Show VPN Peers dialog box 35-22
Show VPNs On Map command 1-33
Show VPNs on Map dialog box 35-21
signatures
adding custom 39-19
cloning 39-21
configuring 39-4
configuring settings 39-30
defining 39-1
detailed information 39-2
editing 39-14
editing Meta engine component list 39-29
editing or tuning parameters 39-23
enabling or disabling 39-14
engines 39-20
exporting 39-9
finding from CS-MARS events 72-45
finding from Event Viewer events 69-54
inheritance 39-3
parameters list 39-24
policy 39-4
selecting category for Cisco IOS IPS 45-6
shortcut menu 39-10
tuning 69-66
tuning recommendations 36-4
understanding 39-1

updates
automatically applying 44-6
checking for and downloading 44-5
configuring server 44-4
managing 44-4
manually applying 44-7
viewing related CS-MARS events 72-44
viewing related events 69-57
viewing update level 39-9, 39-13
Signature Settings page 39-30
Signatures page
overview 39-4
shortcut menu 39-10
Simple Network Management Protocol
see SNMP 49-14
single sign on server (SSO) objects
properties 34-42
SIP (ASA, PIX) class map objects
creating 17-22
SIP (ASA/PIX/FWSM) policy map objects
creating 17-22
properties 17-83, 17-93, 17-102, 17-103
SIP (IOS) class map objects
creating 21-16
match criteria 21-25
SIP (IOS) policy map objects
creating 21-16
match conditions and actions 21-35
SIP class map objects
match criteria 17-85, 17-95
SIP policy map objects
match conditions and actions 17-85, 17-95
Site-to-Site VPN Manager
assigning shared policies 5-44
copying shared policies 5-47
managing policies 5-30
modifying policy assignments 5-49
policy banner 5-38
policy shortcut menu 5-40
renaming policies 5-48
sharing local policies 5-41
unassigning policies 5-36
understanding shared policies 5-37
unsharing policies 5-43
Site-to-Site VPN Manager window 25-18
Site-to-Site VPN policy page (Device view) 25-19
site-to-site VPNs
accessing topologies and policies 25-17
configuring global settings
configuring fragmentation settings 26-31, 26-44
configuring IKEv2 settings 26-37
configuring ISAKMP/IPsec settings 26-33
configuring NAT settings 26-42
overview 26-30
understanding NAT settings 26-41
configuring IKE and IPsec policies 26-1
creating or editing Extranet VPN topologies 25-66
creating or editing VPN topologies 25-28
discovering 25-24
managing 25-1
rediscovering 25-27
repairing discovered VPNs with multiple spoke definitions 25-26
understanding discovery 25-20
understanding topologies 25-2
using device overrides to customize VPN policies 25-13
viewing summary of VPN configuration 25-63
Site-to-Site VPNs command 1-34
Skinny policy map objects
creating 17-22
match conditions and actions 17-89
properties 17-87
SLA monitor objects
attributes 51-10
configuring 51-9
understanding 51-8
Smartfilter (N2H2)

configuring for web filter rules policies 18-15, 18-19

configuring for zone based firewall rules policies 21-36, 21-39, 21-41

Smart Tunnel Auto Signon Entry dialog box 34-72

Smart Tunnel Auto Signon Lists dialog box 34-71

smart tunnels

configuring for ASA SSL VPNs 31-85

SMTP

configuring for inspection rules 17-20

preventing DoS attacks using zone based firewall 21-27

preventing spam using zone based firewall 21-27

SMTP class map objects

creating 21-16

match criteria 21-27

SMTP policy map objects

creating 21-16

match conditions and actions 21-35

SMTP server

configuring 1-27

PIX/ASA/FWSM 52-22

SNEXT 62-46

SNMP

about SNMP Version 3 49-15

Cisco IOS routers

defining agent properties 63-67

enabling traps 63-68

overview 63-66

Permission dialog box 63-70

SNMP Policy page 63-69

SNMP Traps dialog box 63-72

Trap Receiver dialog box 63-71

configuring for HPM S2S polling 71-40

configuring for IPS sensors 36-8

configuring on firewall devices 49-14

IPS general options 36-10

IPS trap options 36-11, 36-13

PIX/ASA/FWSM 49-17

groups 49-24

host access 49-22

MIBs 49-14

OIDs 49-14

SNMPv3 49-24, 49-25

Trap configuration 49-19

users 49-25

terminology 49-15

SNMP Credentials dialog box 3-48

SNMP policy map objects

creating 17-22

properties 17-90

SNMP Trap Communication dialog box 36-12, 36-14

SNMP Trap Communication tab, SNMP policy for IPS 36-11, 36-13

socket read timeout

device communication 11-22

Software Application Support contracts 10-16

Source Contents dialog box 12-14

spam

blocking spam using zone-based firewall rules 21-27

spoke-to-spoke connections, DMVPN 27-10

spoofing, preventing 57-1, 57-3

spoofing attacks, preventing 17-4

SSH

configuring on IOS routers, Catalyst switches, Catalyst 6500/7600 devices 2-6

line ending conventions 2-5

preventing non-SSH connections 2-7

setting up 2-5

testing authentication 2-6

troubleshooting connections 9-7

SSL

remote access SSL VPNs

advanced settings (ASA) 31-72

AnyConnect client settings (ASA) 31-62, 31-64

browser plug-ins 31-60

content rewrite rules (ASA) 31-53

encoding rules (ASA) 31-55

Kerberos Constrained Delegation (KCD on ASA) 31-66, 31-69

proxy bypass rules (ASA) 31-59
remote access VPNs 31-43
access modes 30-4
access policies (ASA), configuring 31-49
access policies (ASA), reference 31-45
access policies (ASA), understanding 31-44
AnyConnect client image settings (ASA) 31-65
AnyConnect custom attributes (ASA) 31-70, 31-71
cluster load balancing 31-5
configuring HTTP/HTTPS proxies and proxy bypass(ASA) 31-57
connection profiles 31-7
connection profiles (ASA) 31-8
Context Editor dialog box (IOS) 33-15, 33-16
creating on ASA 30-14
creating on IOS devices 30-32
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
example 30-3
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
limitations 30-8
managing support files 30-5
NAT settings 26-42
other settings (ASA) 31-51
performance settings (ASA) 31-52
policies (IOS) 33-14
prerequisites 30-7
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9
server certificate verification (ASA) 31-30, 31-32, 31-73
shared license clients (ASA) 31-76
shared licenses (ASA) 31-74
shared license servers (ASA) 31-77
understanding 30-2
understanding NAT settings 26-41
wizard 30-13
setting up 2-3
troubleshooting certificate errors 9-5
VPN
sharing connection profiles on ASAs 30-8
SSL authentication certificates
adding thumbprints manually 9-5
configuring default settings for how handled 11-22
SSL VPN
policy discovery restriction 3-8
SSL VPN Access page (ASA) 31-45
SSL VPN bookmark objects
configuring 31-82
post URL method and macro substitutions 31-84
SSL VPN Bookmarks objects
SSL VPN Bookmarks dialog box 34-45
SSL VPN Configuration wizard
Access page (ASA) 30-16
Connection Profile page (ASA) 30-17
Gateway and Context Page (IOS) 30-33
Portal Page Customization Page (IOS) 30-35
SSL VPN Customization objects
configuring 31-78
creating custom Logon page 31-82
localizing 31-80
SSL VPN gateway objects
properties 34-64
SSL VPN Other Settings page (ASA)
Advanced tab 31-72
Client Settings tab 31-64
Content Rewrite tab 31-53
Encoding tab 31-55
Microsoft KCD Server tab 31-66, 31-69
overview 31-51
Performance tab 31-52

Proxy tab 31-57
SSL Server Verification tab 31-30, 31-32, 31-73
SSL VPN Policy page (IOS) 33-14
SSL VPNs
ASA devices
configuring bookmarks 31-82
configuring portal appearance 31-78
configuring WINS servers for file system access 31-88
customizing 31-77
post URL method and macro substitutions in bookmarks 31-84
smart tunnels 31-85
IOS devices
configuring bookmarks 31-82
configuring WINS servers for file system access 31-88
SSL VPN Shared License page (ASA) 31-74
SSL VPN smart tunnel auto sign-on list objects
attributes 34-71
SSL VPN smart tunnel list objects
attributes 34-66, 34-69
configuring 31-85
stateful failover 50-3, 50-4
site-to-site VPN 25-54
stateless failover 50-3
states
activity 4-4
ticket 4-4
static crypto maps 26-19
Static Group tab (IGMP) 55-6
static NAT
Cisco IOS routers
disable automatic aliasing 24-7
disable payload 24-10
on Cisco IOS routers 24-6
static null 0 routing 56-128
static routes
configuring on firewall devices 56-128
PIX/ASA/FWSM
configuration 56-130
static routing
Cisco IOS routers
defining on 67-50
overview 67-50
Static Routing dialog box 67-52
Static Routing Policy page 67-51
Static Rule
PIX/ASA/FWSM 24-26
add/edit 24-27
status
activity 4-4
ticket 4-4
subinterfaces 46-7, 46-15
specifying during policy definition 6-76
Submit Activity command 1-36
Submit Activity dialog box 4-20
Submit and Deploy command 1-30
Submit command 1-30
Submit Deployment Job dialog box 8-38
Submitted activity state 4-5
Submit Ticket command 1-37
Sun RPC class map objects
creating 21-16
match criteria 21-29
Sun RPC policy map objects
creating 21-16
match conditions and actions 21-35
support, technical
creating diagnostic file 10-28
generating data 10-28
generating deployment or discovery status reports 10-30
generating partial database backup 10-30
Suspend Deployment Schedule dialog box 8-20, 8-58
switches
communication requirements 2-1
SYN flooding attacks, preventing 17-5
syslog

access rule look-up 72-17
deeply parsed for Event Viewer 69-6
logging
PIX/ASA/FWSM 54-1
message properties 69-18
syslog messages supported for policy lookup 72-46
syslog relay
CPU throttling policy 11-33
syslogs
Cisco IOS routers 65-1
system variables
devices 7-7
firewall 7-9
FlexConfigs 7-7
remote access VPN 7-19
routers 7-13
VPN 7-14

T

tables
using 1-48
tables, rules
adding rules 12-9
columns and headings 1-49
commands, Edit menu 1-31
converting IPv4 rules 12-28
cut, copy, and paste rules 12-9
disabling rules 12-20
enabling rules 12-20
filtering 1-48
finding and replacing items 12-16
removing rules 12-9
sections 12-20
using 12-8
TACACS+
description 6-28
settings in AAA server objects 6-38
Take Over User Session page 11-71
Target Value Rating dialog box 40-19
Target Value Ratings, IPS Network Information policy 40-17
target value ratings (IPS) 40-17
task flow
deployment
non-Workflow mode 8-3
Workflow mode 8-5
taskflow 1-19
TCP Map objects
properties 58-22
TCP State Bypass
ASA/FWSM 58-3
Telnet
PIX/ASA/FWSM 49-29
configuration 49-29
text fields
ASCII limitations 1-50
finding text in multiple-line 1-50
navigating 1-50
using 1-49
text objects
creating 7-32
TFTP servers
PIX/ASA/FWSM 52-23
thin client access mode 30-4
thresholds
configuring anomaly detection 41-11
understanding anomaly detection 41-9
throughput
VPN user reports 70-16
ticketing
overview 1-20
Ticket Management
settings 11-72
ticket management
comparing workflow modes 1-22
Ticket Manager window 4-10
tickets

closing 4-16
creating 4-14
discarding 4-22
multiple users 4-4
opening 4-15
states 4-4
Ticket Manager window 4-10
understanding 4-1
using global search to find specific tickets 1-42
validating 4-18
viewing change reports 4-16
viewing status and history 4-23
working with 4-7
Tickets menu 1-36
tiered hub-and-spoke topologies 25-5
time
changing range for reports 70-22
timeouts
on firewall devices 57-4
timeouts (NAT)
Cisco IOS routers 24-13
Timeout Value
Firewall AAA 15-30
time range objects
attributes for recurring ranges 6-72
configuring 6-71
time slider (Event Viewer)
filtering with 69-44
using 69-25
time synchronization
on IOS routers 63-96
time zone settings
certificate errors 9-7
Cisco IOS routers
Clock Policy page 63-23
defining time zone and DST 63-22
overview 63-22
TMS
deploying configurations 8-43
deployment method 8-10
Token Management page 11-73
Token Management System (TMS)
settings 11-73
toolbar
activities 4-8, 4-9
toolbar reference
Configuration Manager 1-39
event table in Event Viewer 69-16
toolbars
Report Manager generated report 70-12
Report Manager report settings 70-10
Tools menu
Configuration Manager 1-34
Report Manager 70-8
Trace Route 72-26
TraceRoute 72-28
traffic class
PIX/ASA/FWSM
rules wizard 58-7
Traffic Classification dialog box 19-12
Traffic Classification tab 19-11
traffic encryption key (KEK), GET VPN 29-4
traffic flow notifications
configuring for IPS 36-30
traffic flow objects
default inspection traffic 58-20
properties 58-18
traffic match criteria 58-2
traffic zones 22-1
asymmetric routing 22-1
benefits 22-1
clustering 22-8
configuring 22-9
Equal-Cost Multi-Path (ECMP) 22-4
failover 22-8
firewall mode 22-8
guidelines 22-8
load balancing 22-1

lost route 22-1
prerequisites 22-7
security levels 22-6
supported services 22-6
understanding 22-6
transactional commit model 60-2
Transactional Commit page
PIX/ASA/FWSM Platform 60-2
transcripts
viewing 8-59
Transcript Viewer window 8-62
transform sets
attributes 26-27
understanding 26-20
Translation Exemption (NAT-0 ACL) Rule
PIX/ASA/FWSM 24-20
add/edit 24-21
Translation Options
PIX/ASA/FWSM 24-16, 24-17
Translation Rules
Add/Edit Per-Session NAT rules dialog boxes 24-47
ASA 8.3+ 24-34
Add/Edit NAT rules dialog boxes 24-36
per-session NAT rules 24-46
PIX/ASA/FWSM 24-19
translation table
clearing on deployment 60-1
transparent bridging
Cisco IOS routers
BVI interfaces 63-19
overview 63-18
defining bridge groups 63-20
transparent firewall
configuring on PIX/ASA/FWSM 47-1
NAT 24-16
transparent rules
adding or editing a rule 23-5
configuring 23-1
configuring DHCP passthrough for IOS devices 23-3
configuring in Map view 35-23
deleting 12-9
disabling 12-20
editing 12-10
editing the EtherType 23-7
editing the mask 23-7
enabling 12-20
managing 23-1
moving 12-19
Transparent Rules page 23-3
understanding processing order 12-2
Transparent Rules page 23-3
transport protocols
device defaults 11-22
overview of device requirements 2-1
transport settings
AUS 2-8
Configuration Engine 2-8
SSH 2-5
SSL (HTTPS) 2-3
traps, SNMP
configuring for IPS sensors 36-8
IPS options 36-11, 36-13
trees
selecting items 1-45
Trend class map objects
creating 21-36
Trend parameter map objects
creating 21-36
properties 21-42
troubleshooting
AUS deployment 9-18
Catalyst switch and module deployment 9-16
Configuration Engine deployment 9-18
creating diagnostics file 10-28
CS-MARS queries 72-40
deleted FWSM contexts do not remove configuration files 59-7
deployment 9-9

device communication and deployment 9-1
device discovery failures 3-7
device managers 72-16
device managers, using 72-14
devices marked with red X in device selector 9-9
error attempting to remove unreferenced object 9-12
Event Manager service status 69-31
Event Viewer Unavailable message 11-27, 11-36, 69-30
FlexConfigs 7-38
FWSM multiple-context deployment failures 9-17
generating data for TAC 10-28
generating deployment or discovery status reports 10-30
GET VPN registration failure 29-9
global correlation (IPS) configuration 42-4
ignoring device errors during deployment 9-10
invalid certificate error 9-7
minimum memory errors for ASA 8.3+ 9-12
mixing deployment methods 9-14
Not able to connect to server message, Report Manager 70-36
online help, problems accessing 1-52
packet capture, using 72-30
packet tracer, using 72-23
policy objects not available in Event Viewer 69-68
preshared key policies in VPN not discovered 25-23
Report Manager 70-36
router connection failures 2-2
router deployment 9-14
Security Manager cannot contact device after deployment 9-12
SSL certificate errors 9-5, 9-6
user interface problems 1-51
VPN crypto traffic unexpectedly dropped on GET VPN interfaces 29-9
VPNs with routing processes 9-13
VRF-aware IPsec deployment failures on Catalyst 6500/7600 devices 25-17
trunk ports
Create and Edit Interface dialog boxes-Trunk Port mode 68-14
understanding 68-5
Trusted Transitive Introduction (TTI)
use in SDP policies 63-81
TrustSec
Add/Edit Connection Peer dialog box 14-13
configuring connection peers 14-13
configuring ISE settings 11-56
configuring SXP 14-8
configuring SXP connection peers 14-12
security group objects
creating 14-14
SGT role mapping 14-11
TrustSec firewall policies
configuring 14-7
configuring rules 14-17
managing 14-1
TrustSec policies
monitoring 14-17
TrustSec security group objects
selecting 14-16

U

Unassign Policy command 1-32
Undock Map View command 1-33
unicast rekey in GET VPN 29-6
Unicast Reverse Path Forwarding 57-1, 57-3
unicast reverse path forwarding
enabling on routers 62-20
Unshare Policy command 1-32
Unspecified Bit Rate (UBR) 62-49
Unspecified Bit Rate Plus (UBR+) 62-49
Update Level dialog box 39-9, 39-13
updating images on devices 73-20
Updating Licenses from File dialog box 11-61
Updating Licenses via CCO dialog box 11-60
URLF Glob parameter map objects

metacharacters 21-46
properties 21-45
URL Filter parameter map objects
creating 21-36
properties 21-43
usage reports
generating 6-15
user accounts
configuring IPS 36-18
configuring IPS password requirements 36-20
discovery and deployment of IPS 36-17
IPS account attributes 36-19
managing IPS device 36-15
PIX/ASA/FWSM 51-7
add/edit 51-7
rolling back configurations 8-64
understanding IPS user roles 36-15
understanding managed and unmanaged passwords 36-16
User Accounts policy, IPS devices 36-18
user group objects
advanced PIX 6.3 settings 34-82
browser proxy settings 34-87
clientless settings 34-83
client VPN software update (IOS) settings 34-81
DNS/WINS settings 34-77
general settings 34-75
IOS client settings 34-78
IOS Xauth settings 34-80
split tunneling settings (Easy VPN/remote access IPSec VPN) 34-77
SSL VPN connection settings 34-88
SSL VPN full tunnel settings 34-84
SSL VPN split tunneling settings 34-86
technology settings 34-73
thin client settings 34-84
user group policies
configuring for Easy VPN 28-14
configuring for remote access IPsec VPNs on IOS/PIX 6.3 33-13
User Group Policy page 33-13
user identity acquisition 13-2
user interface
applications overview 1-6
basic features 1-29
dialog box too big for screen 1-52
freezing 1-51
how permissions affect what you can do 1-11
Java errors 1-52
maps toolbar reference 35-4
map view 35-1
menu reference for Configuration Manager 1-29
missing text 1-52
overview of Configuration Manager 1-14
rules tables 12-8
searching for items 1-42
selecting items in a tree 1-45
selecting or specifying files 1-50
table
columns and headings 1-49
sections 12-20
tables 1-48
text fields
ASCII limitations 1-50
finding text in multiple-line 1-50
navigating 1-50
using 1-49
toolbars
Configuration Manager 1-39
event table in Event Viewer 69-16
troubleshooting 1-51
wizards 1-47
user login credentials for device access 3-4
user passwords
changing 10-24
user preferences
PIX/ASA/FWSM 60-1
Deployment page 60-1
Transactional Commit page 60-2

user roles, IPS 36-15
users
how permissions affect what you can do 1-11
taking over configuration session 10-23
User Statistics
MPC rule wizard
tab 58-8
user statistics, collecting 13-25
user taskflow 1-19

V

Validate Activity command 1-36
Validate command 1-30
Validate Ticket command 1-37
Validation dialog box 4-18
validation error messages 4-18
Values Assignment dialog box 7-37
Variable Bit Rate-Non-Real Time (VBR-nrt) 62-49
Variable Bit Rate-Real Time (VBR-rt) 62-49
variables
deleting FlexConfig 7-28
FlexConfig objects 7-5, 7-6
changing variable values 7-35
VDI servers 34-15
Velocity Engine error message 7-38
Velocity Template Engine
scripting language 7-3
View Changes command 1-30, 1-36, 1-37
viewing interface allocations 59-12
View menu
Configuration Manager 1-31
Event Viewer 69-10
views
Device 1-15
Event Viewer
clearing filters 69-48
column based filters 69-45
event based filters 69-47
filtering overview 69-43
refreshing event table 69-44
selecting time range 69-43
switching between real-time and historical 69-42
text searches (quick filter) 69-47
using time slider with filtering 69-44
HPM 71-21
column-based filters 71-17
Map 1-18
overview 1-14
Policy 1-16
views (Event Viewer)
arranging 69-38
configuring color rules 69-40
creating custom 69-41
customizing event table appearance 69-39
deleting custom 69-43
editing description 69-41
editing name 69-41
Event Monitoring window overview 69-14
Event Viewer overview 69-7
floating 69-38
list 69-12
opening 69-38
overview 69-3
saving 69-42
using 69-37
virtual ASA
about 46-1
virtual channel identifier (VCI) 62-47
virtual firewalls
See security contexts
virtual fragment reassembly (VFR) 62-19
virtual path identifier (VPI) 62-47
Virtual Routing Forwarding (VRF)
VRF-Aware IPsec 25-14
virtual sensors
advantages 38-3
assigning interfaces 38-4

attributes 38-7
configuring 38-1, 38-5
deleting 38-10
discovering policies 5-14
editing policies 38-9
identifying 38-5
inline TCP session tracking mode 38-3
Normalizer mode 38-4
renaming 38-8
restrictions 38-3
showing containment 3-56
understanding 38-1
Virtual Sensors page 38-5
virtual terminal (VTY)
Cisco IOS routers
defining AAA settings 63-40
defining line groups 63-38
defining line setup parameters 63-38
virtual terminal (VTY) lines
Cisco IOS routers
VTY Line dialog box 63-51
VTY Policy page 63-50
VLAN
configuring IPS groups 37-15
configuring IPS inline pairs 37-14
VLAN ACLs (VACLs)
defining 68-37
deleting 68-38
understanding 68-36
VLAN access maps 68-37
VLANs
Catalyst switches and 7600 Series routers
Create and Edit VLAN ACL Content dialog boxes 68-41
Create and Edit VLAN ACL dialog boxes 68-41
Create and Edit VLAN dialog boxes 68-28
defining 68-26
defining Data Port for IDSM 68-46
defining EtherChannel for IDSM 68-44
defining groups 68-32
defining VACLs 68-37
deleting 68-27
deleting Data Port for IDSM 68-47
deleting EtherChannel for IDSM 68-45
deleting groups 68-33
deleting VACLs 68-38
Interfaces/VLANs page-VLANs tab 68-27
understanding 68-25
understanding VACLs 68-36
understanding VLAN groups 68-31
VLAN Access Lists page 68-39
VPDN groups 46-72
VPN
configuring policy defaults 11-74, 25-12
mixing deployment methods 9-14
policy discovery restriction for web VPNs 3-8
Report Manager reports
general VPN reports 70-16
VPN top reports 70-16
system variables 7-14
traffic sent unencrypted 9-15
updating routing processes 9-13
using device overrides to customize VPN policies 25-13
zone-based firewall 21-6
VPN default policies
configuring 25-12
factory defaults 25-12
understanding 25-12
VPN discovery
prerequisites 25-21
procedure 25-24
rules 25-22
supported and unsupported technologies and topologies 25-20
understanding 25-20
VPN global settings
GET VPN

VPN Global Settings for GET page 29-16
VPN Global Settings policy
General Settings tab 26-31, 26-44
IKEv2 tab 26-37
ISAKMP/IPsec tab 26-33
NAT Settings tab 26-42
VPN Peers dialog box 35-22
VPN Policy Defaults page 11-74
VPN rediscovery 25-27
VPNs
AAA services 48-4
ASA devices
configuring bookmarks 31-82
configuring portal appearance 31-78
configuring WINS servers for file system access 31-88
customizing 31-77
post URL method and macro substitutions in bookmarks 31-84
smart tunnels 31-85
configuring remote access using wizard 30-13
creating in Map view 35-21
Easy VPN
connection profiles 28-13
connection profiles (ASA, PIX 7+) 31-8
IOS devices
configuring bookmarks 31-82
configuring WINS servers for file system access 31-88
IPsec
access policies for IKEv2 (ASA), configuring 31-49
access policies for IKEv2 (ASA), reference 31-45
access policies for IKEv2 (ASA), understanding 31-44
certificate to connection profile map policy (IKEv1) 31-36
certificate to connection profile map rules (IKEv1) 31-37
cluster load balancing 31-5
configuring IKE and IPsec policies 26-1
connection profiles 31-7
connection profiles (ASA, PIX 7+) 31-8
creating on ASA/PIX 7.0+ 30-25
creating on IOS/PIX 6.3+ 30-36
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
Dynamic VTI/VRF Aware IPsec settings 33-7
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
high availability policies 33-11
IKE proposals 26-9
IKEv2 authentication 26-68, 26-70, 26-72
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
NAT settings 26-42
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9
understanding IKE 26-5
understanding NAT settings 26-41
user group policies for IOS, PIX 6.3 33-13
VPNSM, VPN SPA, VSPA settings 33-6
IPsec proposals
attributes for ASA and PIX 7.0+ devices 31-41
attributes for IOS and PIX 6.3 devices 33-4
configuring for ASA and PIX 7.0+ devices 31-40
configuring for IOS and PIX 6.3 devices 33-3
Map view 35-20
policy discovery 5-12
remote access
access modes 30-4
device support 30-8
discovering 30-12
managing 30-1
managing (ASA, PIX 7.0+) 31-1
managing (IOS, PIX 6.3) 33-1

SSL 31-43
remote access IPSec
understanding 30-2
remote access SSL
example 30-3
limitations 30-8
managing support files 30-5
prerequisites 30-7
understanding 30-2
shared policies 5-4
site-to-site
configuring IKE and IPsec policies 26-1
policies overview 25-8
site-to-site VPNs 25-1
SSL
access policies (ASA), configuring 31-49
access policies (ASA), reference 31-45
access policies (ASA), understanding 31-44
advanced settings (ASA) 31-72
AnyConnect client image settings (ASA) 31-65
AnyConnect client settings (ASA) 31-62, 31-64
AnyConnect custom attributes (ASA) 31-70, 31-71
browser plug-ins (ASA) 31-60
cluster load balancing 31-5
configuring HTTP/HTTPS proxies and proxy bypass(ASA) 31-57
connection profiles 31-7
connection profiles (ASA) 31-8
content rewrite rules (ASA) 31-53
Context Editor dialog box (IOS) 33-15, 33-16
creating on ASA 30-14
creating on IOS devices 30-32
dynamic access policies 32-1, 32-2
dynamic access policy (DAP) attributes 32-4, 32-7
Dynamic Access policy page (ASA) 32-11
encoding rules (ASA) 31-55
fragmentation settings 26-31, 26-44
global settings 26-30
group policies, configuring 31-26
group policies, creating 31-28
group policies, understanding 31-27
IKEv2 settings 26-37
ISAKMP/IPsec settings 26-33
Kerberos Constrained Delegation (KCD on ASA) 31-66, 31-69
NAT settings 26-42
other settings (ASA) 31-51
performance settings (ASA) 31-52
policies (IOS) 33-14
proxy bypass rules (ASA) 31-59
public key infrastructure (PKI) policies 26-56
secure desktop manager policies 32-9
server certificate verification (ASA) 31-30, 31-32, 31-73
shared license (ASA) 31-74
shared license clients (ASA) 31-76
shared license servers (ASA) 31-77
understanding NAT settings 26-41
wizard 30-13
understanding 30-1
VPN Service Port Adapters (VSPAs)
configuring 25-42
VPN Services Module (VPNSM)
configuring 25-42
VPN Shared Port Adapter (VPN SPA)
configuring 25-42
VPNSM/VPN SPA/VSPA Settings dialog box 33-6
VPN Summary page 25-63
VPN topologies
accessing 25-17
assigning initial policies to new 25-62
assigning shared policies 5-44
cloning device VPN assignments 3-56
cloning shared policies 5-47
configuring dial backup 25-40
configuring GET VPN peers 25-60
configuring in Device view 25-19
creating or editing 25-28

creating or editing Extranet 25-66
defining endpoints and protected networks 25-34
defining GET VPN group encryption 25-54
deleting 25-71
discovering 25-20, 25-24
full mesh 25-4
hub-and-spoke 25-2
including unmanaged or non-Cisco devices 25-11
joined hub-and-spoke 25-5
locking 5-10
naming 25-30
partial mesh 25-5
point-to-point 25-3
rediscovering 25-27
removing devices 25-32
renaming policies 5-48
repairing discovered VPNs with multiple spoke definitions 25-26
selecting devices 25-32
tiered hub-and-spoke 25-5
unassigning policies 5-36
understanding 25-2
unsharing policies 5-43
using device overrides to customize VPN policies 25-13
viewing summary of VPN configuration 25-63
VRF-Aware IPsec
changing on Catalyst switches and 7600 routers 25-17
configuring 25-48
one-box solution 25-14
two-box solution 25-15
understanding 25-14
VRF-Aware IPsec tab (site-to-site VPN) 25-48
VTP modes, for Catalyst switches 68-1
VTY Line dialog box 63-51
Accounting tab 63-57
Authentication tab 63-55
Authorization tab 63-56
Setup tab 63-52

W

WAN interface card (WIC) 62-36
Warning - Partial VPN Deployment dialog box 8-31
warnings
significance of 2-lxiii
Web Filter policy map objects
creating 21-36
match conditions and actions 21-35
properties 21-47
web filter rules
ACL naming conventions 12-5
ASA/FWSM/PIX
converting IPv4 12-28
deleting 12-9
editing 12-10
moving 12-19
attributes (IOS) 18-13
configuring exclusive domains for IOS devices 18-10
configuring for ASA, PIX, FWSM devices 18-2
configuring for IOS devices 18-10
configuring in Map view 35-23
disabling 12-20
enabling 12-20
exclusive domain names (IOS) 18-14
managing 18-1
preserving ACL names 12-4
understanding 18-1
understanding NAT effects 12-3
understanding processing order 12-2
Web Filter Rules page (ASA/FWSM/PIX) 18-3
Web Filter Rules page (IOS) 18-12
web filter server properties 18-19
Web Filter Rules page (ASA/FWSM/PIX) 18-3
Web Filter Rules page (IOS) 18-12
Web Filter Server Configuration dialog box 18-19
web filter servers
attributes 18-19
configuring settings 18-15

configuring settings in Map view 35-24

configuring zone-based firewall settings in Map view 35-24

Web Filter settings page 18-16

Websense

configuring for web filter rules policies 18-15, 18-19

configuring for zone based firewall rules policies 21-36, 21-39, 21-41

Websense class map objects

creating 21-36

match criteria 21-30

Websense parameter map objects

creating 21-36

properties 21-39

web VPN

policy discovery restriction 3-8

Weighted Random Early Detection (WRED) 66-4

Whitelist/Blacklist tab 19-14

windows

arranging report 70-30

arranging views 69-38

closing report 70-32

undocking maps 35-2

Windows Messenger class map objects

creating 21-16

match criteria 21-21

Windows NT servers

use by ASA, PIX, and FWSM devices 6-29

Windows Server 2012 security settings 10-2

WINS Server Lists objects

attributes 34-90

creating 31-88

wizard

installation manager 73-26

wizards

configuring remote access SSL VPNs on ASA devices 30-14

configuring remote access SSL VPNs on IOS devices 30-32

configuring remote access VPNs 30-13

Copy Policies 5-33

Create Extranet VPN Topology 25-66

Create VPN Topology 25-28

creating remote access IPsec VPNs on ASA/PIX 7.0+ devices 30-25

creating remote access IPsec VPNs on IOS/PIX 6.3 devices 30-36

creating user group policies 30-20

Discover VPN policies 25-24

New Device 3-6

Rediscover VPN policies 25-27

Share Policies 5-42

wizards, using 1-47

workflow

overview 1-20

Workflow mode

changing modes 1-28

comparing with non-Workflow mode 1-22

configuration files

deploying 8-34, 8-39

previewing 8-44

configurations

rolling back 8-69

creating activities 4-14

deployment

viewing device details 8-26

viewing job history 8-26

jobs

aborting 8-55

approving 8-39

discarding 8-41

rejecting 8-39

states 8-6

submitting 8-38

opening activities 4-15

understanding 1-21

workflow modes

changing 1-28

comparing 1-22

Workflow Settings page 11-75
working with 3-59
worms
configuring IPS anomaly detection signatures 41-4
understanding 41-2
understanding IPS anomaly detection 41-1
understanding when to turn off anomaly detection 41-4

X

xdm-launcher.exe
device manager 72-16
XLATE table
clearing on deployment 60-1

Y

Yahoo Messenger class map objects
creating 21-16
match criteria 21-21

Z

zone-based firewall
add/edit zones 21-53
advanced options 21-67
changing the default drop rule 21-48
configuring PAM 21-69
configuring rules 21-13, 21-62
configuring settings 21-49
configuring settings in Map view 35-24
Content Filter tab 21-52
designing network zones 21-1
development overview 21-12
general recommendations 21-12
Global Parameters tab 21-50
IPSec VPN 21-6
logging 21-1
overview 21-1
page 21-50
preserving ACL names 12-4
protocol selection 21-68
restrictions 21-3
rules table 21-58
Self zone 21-5
tabs 21-49
troubleshooting 21-54
understanding 21-3
understanding NAT effects 12-3
understanding permit/deny and action 21-8
understanding processing order 12-2
understanding services and protocols 21-11
VPN tab 21-50
VRF 21-7
WAAS tab 21-50
Zones tab 21-50
zone-based firewall rules
configuring in Map view 35-23
deleting 12-9
disabling 12-20
editing 12-10
enabling 12-20
moving 12-19
zone-based firewall rules policies
blocking spam using zone-based firewall rules 21-27
configuring map objects for content filtering rules 21-36
configuring map objects for inspection rules 21-16
creating zones 6-74
inspection parameters 21-31
match conditions for IM applications 21-21
match conditions for P2P applications 21-21
preventing SMTP DoS attacks 21-27
protocol information for IM application inspection 21-33
understanding interface role objects 6-73
Zone Contents dialog box 12-14

zones
creating 6-74
understanding interface role objects 6-73
zones, anomaly detection 41-3
Zoom In command 1-33
Zoom Out command 1-33